

Consistency of restricted maximum likelihood estimators of principal components

Running Title: Consistency of REML estimators

Debashis Paul¹ and Jie Peng

Department of Statistics, University of California, Davis

Abstract

In this paper we consider two closely related problems : estimation of eigenvalues and eigenfunctions of the covariance kernel of functional data based on (possibly) irregular measurements, and the problem of estimating the eigenvalues and eigenvectors of the covariance matrix for high-dimensional Gaussian vectors. In Peng and Paul (2007), a restricted maximum likelihood (REML) approach has been developed to deal with the first problem. In this paper, we establish consistency and derive rate of convergence of the REML estimator for the functional data case, under appropriate smoothness conditions. Moreover, we prove that when the number of measurements per sample curve is bounded, under squared-error loss, the rate of convergence of the REML estimators of eigenfunctions is near-optimal. In the case of Gaussian vectors, asymptotic consistency and an efficient score representation of the estimators are obtained under the assumption that the effective dimension grows at a rate slower than the sample size. These results are derived through an explicit utilization of the intrinsic geometry of the parameter space, which is non-Euclidean. Moreover, the results derived in this paper suggest an asymptotic equivalence between the inference on functional data with dense measurements and that of the high dimensional Gaussian vectors.

1 Introduction

Analysis of functional data, where the measurements per subject, or replicate, are taken on a finite interval, has been one of the growing branches of statistics in recent times. In fields such as longitudinal data analysis, chemometrics, econometrics, the functional data analysis viewpoint has been successfully used to summarize data and gain better understanding of the problems at hand. The monographs of Ramsay and Silverman (2005) and Ferraty and Vieu (2006) give detailed accounts of the applications of functional data approach to various problems in these fields. Depending on how the individual curves are measured, one can think of two different scenarios - (i) when the curves are measured on a dense grid; and (ii) when the measurements are observed on an irregular, and typically sparse set of points on an interval. The first situation usually arises when the data are recorded by some automated instrument, e.g. in chemometrics, where the curves represent the spectra of certain chemical substances. The second scenario is more typical in longitudinal studies where the individual curves could represent the level of concentration of some substance, and the measurements on the subjects may be taken only at irregular time points. In the first scenario, i.e., data on a regular grid, as long as the individual curves are smooth, the measurement noise level is low, and the grid is dense enough, one can essentially treat the data to be on a continuum, and employ techniques similar to the ones used in classical multivariate

¹Paul's research was partially supported by the National Science Foundation under Agreement No. DMS-0112069, as part of a research fellowship from SAMSI.

AMS 2000 subject classifications. Primary 62G20; Secondary 62H25.

Key words and phrases. Functional data analysis, principal component analysis, high-dimensional data, Stiefel manifold, intrinsic geometry, consistency

analysis. For example, Hall and Hosseini-Nasab (2006) derive stochastic expansions of sample PCA when the sample curves are noise-free and measured on a continuum. However, in the second scenario, the irregular nature of the data, and the presence of measurement noise pose challenges and require a different treatment. Under such a scenario, data corresponding to individual subjects can be viewed as partially observed, and noise-corrupted, independent realizations of an underlying stochastic process. The estimation of the eigenvalues and eigenfunctions of a smooth covariance kernel, from sparse, irregular measurements, has been studied by various authors including James, Hastie and Sugar (2000), Yao, Müller and Wang (2005), and Peng and Paul (2007), among others.

In Peng and Paul (2007), a *restricted maximum likelihood (REML)* approach is taken to obtain the estimators. REML estimators are widely used and studied in statistics. For example, the usefulness of REML and profile REML estimation has been recently demonstrated in the context of functional linear mixed effects model by Antoniadis and Sapatinas (2007). In Peng and Paul (2007), it is assumed that the covariance kernel can be well-approximated by a positive-semidefinite kernel of finite rank r whose eigenfunctions can be represented by $M(\geq r)$ known orthonormal basis functions. Thus the basis coefficient matrix B of the approximant belongs to the *Stiefel manifold* of $M \times r$ matrices with orthonormal columns. The working assumption of Gaussianity allows the authors to derive the log-likelihood of the observed data given the measurement times. Then a Newton-Raphson procedure, that respects the geometry of the parameter space, is employed to obtain the estimates by maximizing the log-likelihood. This procedure is based on the formulation of a general Newton-Raphson scheme on Stiefel manifold developed in Edelman, Arias and Smith (1998). Peng and Paul (2007) also derive a computationally efficient approximate cross-validation score for selecting M and r . Through extensive simulation studies, it is demonstrated that the REML estimator is much more efficient than an alternative procedure (Yao *et al.*, 2005) based on local linear smoothing of empirical covariances. The latter estimator does not naturally reside in the parameter space, even though it has been proved to achieve the optimal non-parametric convergence rate in the minimax sense under l^2 loss, under the optimal choice of the bandwidth and when the number of measurements per curve is bounded (Hall, Müller and Wang, 2006). Also, in most situations, our method outperforms the EM approach of James *et al.* (2000). Although the latter estimator also aims to maximize the log-likelihood, it does not naturally reside in the parameter space either, and thus it does not utilize its geometry efficiently.

The superior numerical performance of the REML estimator motivates us to conduct a detailed study of its asymptotic properties. In this paper, we establish consistency and derive the rate of convergence (under l^2 loss) of the REML estimator when the eigenfunctions have a certain degree of smoothness, and when a stable and smooth basis, e.g., the cubic B-spline basis with a pre-determined set of knots, is used for approximating them. The techniques used to prove consistency differ from the standard asymptotic analysis tools when the parameter space is Euclidean. Specifically, we restrict our attention to small ellipsoids around zero in the tangent space to establish a mathematically manageable neighborhood around an “optimal parameter” (a good approximation of the “true parameter” within the *model space*). We derive asymptotic results when the number of measurements per curve grows sufficiently *slowly* with the sample size (referred as the *sparse case*). We also show that for a special scenario of the *sparse case*, when there is a bounded number of measurements per curve, the risk of the REML estimator (measured in squared-error loss) of the eigenfunctions has asymptotically near-optimal rate (i.e., within a factor of $\log n$ of the optimal rate) under an appropriate choice of the number of basis functions.

Besides the *sparse case*, we consider two other closely related problems: (i) the estimation of the eigenvalues and eigenfunctions of a smooth covariance kernel, from dense, possibly irregular,

measurements (referred as the *dense case*); and (ii) the estimation of the eigenvalues and eigenvectors of a high-dimensional covariance matrix (referred as the *matrix case*). In the *matrix case*, we assume that there is preliminary information so that the data can be efficiently approximated in a lower dimensional *known* linear space whose effective dimension grows at a rate slower than the sample size n . The proofs of the results in all three cases utilize the intrinsic geometry of the parameter space through a decomposition of the Kullback-Leibler divergence. However, the *matrix case* and the *dense case* are more closely related, and the techniques for proving the results in these cases are different in certain aspects from the treatment of the *sparse case*, as described in Sections 2 and 3.

Moreover, in the *matrix case*, we also derive a *semiparametric efficient score representation* of the REML estimator (Theorem 4), that is given in terms of the *intrinsic Fisher information operator* (note that the residual term is not necessarily $o_P(n^{-1/2})$). This result is new, and explicitly quantifies the role of the intrinsic geometry of the parameter space on the asymptotic behavior of the estimators. Subsequently, it points to an *asymptotic optimality* of the REML estimators. Here, asymptotic optimality means achieving the asymptotic minimax risk under l^2 loss within a suitable class of covariance matrices (kernels). We want to point out that, in the *matrix case*, the REML estimators coincide with the usual PCA estimates, i.e., the eigenvalues and eigenvectors of the sample covariance matrix (Muirhead, 1982). In Paul and Johnstone (2007), a first order approximation of the PCA estimators is obtained by matrix perturbation analysis. Our current results show that the efficient score representation coincides with this approximation, and thereby gives a geometric interpretation to this. The asymptotically optimal rate of the l^2 -risk of the REML estimator in the *matrix case* follows from this representation and the lower bound on the minimax rate obtained in Paul and Johnstone (2007). Asymptotic properties of high-dimensional PCA under a similar context have also been studied by Fan, Fan and Lv (2007). Recently several approaches have been proposed for estimating large dimensional covariance matrices and their eigenvalues and eigenvectors under suitable sparsity assumptions on the population covariance, e.g. Bickel and Levina (2007, 2008) and El Karoui (2008).

At this point, we would like to highlight the main contributions of this paper. First, we have established the consistency and derived the rate of convergence of REML estimators for functional principal components in two different regimes - the *sparse case* and the *dense case*. In Hall *et al.* (2006), it is shown that an estimator of functional principal component based on a local polynomial approach achieves the optimal nonparametric rate when the number of measurements per curve is bounded. However, to the best of our knowledge, no results exist regarding the consistency, or rate of convergence, of the REML estimators in the functional data context. Secondly, we have derived an efficient score representation for sample principal components of high-dimensional, i.i.d. Gaussian vectors. This involves calculation of the intrinsic Fisher information operator and its inverse, and along the line we also provide an independent verification that the REML estimates under a rank-restricted covariance model are indeed the PCA estimates. Thirdly, we expect that the current framework can be refined to establish efficient score representation of the REML estimators of the functional principal components, and therefore the results obtained in this paper serve as first steps towards studying the asymptotic optimality of these estimators. Moreover, results obtained in this paper suggest an *asymptotic equivalence* between the inference on functional data with dense measurements and that of the high dimensional Gaussian vectors. Finally, our work provides useful techniques for dealing with the analysis of estimation procedures based on minimization of a loss function (e.g. MLE, or more generally M-estimators) over a non-Euclidean parameter space for semiparametric problems. There has been some work on analysis of maximum likelihood

estimators for parametric problems when the parameter space is non-Euclidean (see e.g. Oller and Corcuera, 1995). However, there has been very limited work for non/semi-parametric problems with non-Euclidean parameter space. Recently, Chen and Bickel (2006) establish semiparametric efficiency of estimators in ICA (Independent Component Analysis) problems using a sieve maximum likelihood approach.

The rest of the paper is organized as follows. In Section 2, we present the data model for the functional principal components, and state the consistency results of the REML estimators. In Section 3, we describe the model for high-dimensional Gaussian vectors and derive asymptotic consistency and an efficient score representation of the corresponding REML estimators. Section 4 is devoted to giving an overview of the proof of the consistency result for the functional data case (Theorems 1 and 2). Section 5 gives an outline of the proof of consistency in the *matrix case* (Theorem 3), in particular emphasizing the major differences with the proof of Theorem 1. Section 6 is concerned with the proof of the score representation in the *matrix case* (Theorem 4). Section 7 has a summary of the results and a discussion on some future works. Technical details are given in the appendices.

2 Functional data

In this section, we start with a description of the functional principal components analysis, and then make a distinction between the *sparse case* and the *dense case*. We then present the asymptotic results and relevant conditions for consistency under these two settings.

2.1 Model

Suppose that we observe data $Y_i = (Y_{ij})_{j=1}^{m_i}$, at the design points $T_i = (T_{ij})_{j=1}^{m_i}$, $i = 1, \dots, n$, with

$$Y_{ij} = X_i(T_{ij}) + \sigma \varepsilon_{ij}, \quad (1)$$

where $\{\varepsilon_{ij}\}$ are i.i.d. $N(0, 1)$, $X_i(\cdot)$ are i.i.d. Gaussian processes on the interval $[0, 1]$ (or, more generally, $[a, b]$ for some $a < b$) with mean 0 and covariance kernel $\bar{\Sigma}_0(u, v) = \mathbb{E}[X_i(u)X_i(v)]$. $\bar{\Sigma}_0$ has the spectral decomposition

$$\bar{\Sigma}_0(u, v) = \sum_{k=1}^{\infty} \bar{\lambda}_k \bar{\psi}_k(u) \bar{\psi}_k(v)$$

where $\{\bar{\psi}_k\}_{k=1}^{\infty}$ are orthonormal eigenfunctions and $\bar{\lambda}_1 > \dots > \bar{\lambda}_r > \bar{\lambda}_{r+1} \geq \dots \geq 0$ are the eigenvalues. The assumption that the stochastic process has mean zero is simply to focus only on the asymptotics of the estimates of eigenvalues and eigenfunctions of the covariance kernel (i.e., the functional principal components).

Throughout this paper we assume Gaussianity of the observations. We want to emphasize that, Gaussianity is more of a working assumption in deriving the REML estimators. But it plays a less significant role in asymptotic analysis. For the functional data case, only place where Gaussianity is used is in the proof of Proposition 3, and even this can be relaxed by assuming appropriate tail behavior of the observations. Gaussianity is more crucial in the analysis for the matrix case. The proofs of Proposition 6 and Theorem 4 depend on an exponential inequality on the extreme eigenvalues of a Wishart matrix (based on a result of Davidson and Szarek (2001)), even though we expect the non-asymptotic bound to hold more generally.

In this paper, we are primarily interested in the situation where the design points are i.i.d. from a distribution with density g (random design). We shall consider two scenarios, to be referred as the *sparse case* and the *dense case*, respectively. The *sparse case* refers to the situation when the number of measurements, m_i , are comparatively small (see **B1**). The *dense case* refers to the situation where the m_i 's are large so that the design matrix (i.e., the matrix of basis functions evaluated at the time points) has a concentration property (see **B1'** and **D**). In the latter case, we also allow for the possibility that the design is non-random.

Next, we describe the *model space*, to be denoted by $\mathcal{M}_{M,r} := \mathcal{M}_{M,r}(\phi)$, (for $1 \leq r \leq M$) for the REML estimation procedure. The model space $\mathcal{M}_{M,r}$ consists of the class of covariance kernels $C(\cdot, \cdot)$, which have rank r , and whose eigenfunctions are represented in a known orthonormal basis $\{\phi_k\}_{k=1}^M$ of smooth functions. Furthermore, the nonzero eigenvalues are all distinct. For example, in Peng and Paul (2007), $\{\phi_k\}_{k=1}^M$ is taken to be an orthonormalized cubic B -spline basis with equally spaced knots. Thus, the model space consists of the elements $C(\cdot, \cdot) = \sum_{k=1}^r \lambda_k \psi_k(\cdot) \psi_k(\cdot)$, where $\lambda_1 > \dots > \lambda_r > 0$, and $(\psi_1(\cdot), \dots, \psi_r(\cdot)) = (\phi(\cdot))^T B$, where B is an $M \times r$ matrix satisfying $B^T B = I_r$, and $\phi(\cdot) = (\phi_1(\cdot), \dots, \phi_M(\cdot))^T$. Note that we do not assume that $\bar{\Sigma}_0$ belongs to the *model space*. For the asymptotic analysis, we only assume that it can be well-approximated by a member of the *model space* (see condition **C** and *Lemma 1*). We define the *best approximation error* of the model as $\inf_{\tilde{C} \in \mathcal{M}_{M,r}(\phi)} \|\bar{\Sigma}_0 - \tilde{C}\|_F$, where $\|\cdot\|_F$ denotes the Hilbert-Schmidt norm. A rank r approximation to $\bar{\Sigma}_0$ in $\mathcal{M}_{M,r}(\phi)$ can be defined as

$$\Sigma_{*0}(u, v) = \sum_{k=1}^r \lambda_{*k} \psi_{*k}(u) \psi_{*k}(v),$$

with $\lambda_{*1} > \dots > \lambda_{*r} > 0$, and

$$(\psi_{*1}(t), \dots, \psi_{*r}(t)) = (\phi(t))^T B_*,$$

where B_* is an $M \times r$ matrix satisfying $B_*^T B_* = I_r$. We refer to $\{(\psi_{*k}, \lambda_{*k})\}_{k=1}^r$, or equivalently, the pair (B_*, Λ_*) , as an *optimal parameter*, if the corresponding Σ_{*0} is a close approximation to $\bar{\Sigma}_0$ in the sense that, the approximation error $\|\bar{\Sigma}_0 - \Sigma_{*0}\|_F$ has the same rate (as a function of M) as the best approximation error. Henceforth, (B_*, Λ_*) is used to denote an optimal parameter.

Observe that, under model (1), Y_i are independent, and conditionally on T_i they are distributed as $N_{m_i}(0, \bar{\Sigma}_i)$. Here, the $m_i \times m_i$ matrix $\bar{\Sigma}_i$ is of the form $\bar{\Sigma}_i = ((\bar{\Sigma}_0(T_{ij}, T_{ij'}))_{j,j'=1}^{m_i}) + \sigma^2 I_{m_i}$. Then the matrix $\Sigma_{*i} = \Phi_i^T B_* \Lambda_* B_*^T \Phi_i + \sigma^2 I_{m_i}$ is an approximation to $\bar{\Sigma}_i$, where $\Phi_i := [\phi(T_{i1}) : \dots : \phi(T_{im_i})]$ is an $M \times m_i$ matrix. We shall use Λ to denote interchangeably the $r \times r$ diagonal matrix $\text{diag}(\lambda_1, \dots, \lambda_r)$ and the $r \times 1$ vector $(\lambda_1, \dots, \lambda_r)^T$. Note that, the parameter (B, Λ) belongs to the parameter space $\Omega := \mathcal{S}_{M,r} \otimes \mathbb{R}_+^r$, where $\mathcal{S}_{M,r} = \{A \in \mathbb{R}^{M \times r} : A^T A = I_r\}$ is the *Stiefel manifold* of $M \times r$ matrices with orthonormal columns. For fixed r and M , the *REML estimator* of $\{(\bar{\psi}_k, \bar{\lambda}_k)\}_{k=1}^r$ is defined as a minimizer over Ω of the negative log-likelihood (up to an additive constant and the scale factor n):

$$L_n(B, \Lambda) = \frac{1}{2n} \sum_{i=1}^n \text{tr}(\Sigma_i^{-1} Y_i Y_i^T) + \frac{1}{2n} \sum_{i=1}^n \log |\Sigma_i|, \quad (2)$$

where $\Sigma_i = \Phi_i^T B \Lambda B^T \Phi_i + \sigma^2 I_{m_i}$.

2.2 Consistency

We shall present results on consistency of the REML estimators of functional principal components in the two different regimes considered above, namely, the *sparse case* (i.e., when the number of measurements per curve is “small”) and the *dense case* (i.e., when the number of measurements per curve is “large”). Throughout this paper, we assume that σ^2 is known, even though Peng and Paul (2007) provide estimate of σ^2 as well. This assumption is primarily to simplify the exposition. It can be verified that all the consistency results derived in this paper hold even when σ^2 is estimated. We make the following assumptions about the covariance kernel $\bar{\Sigma}_0$.

- A1** The r largest eigenvalues of $\bar{\Sigma}_0$ satisfy, (i) $c_1 \geq \bar{\lambda}_1 > \dots > \bar{\lambda}_r > \bar{\lambda}_{r+1}$ for some $c_1 < \infty$; (ii) $\max_{1 \leq j \leq r} (\bar{\lambda}_j - \bar{\lambda}_{j+1})^{-1} \leq c_2 < \infty$.
- A2** The eigenfunctions $\{\bar{\psi}_k\}_{k=1}^r$ are four times continuously differentiable and satisfy

$$\max_{1 \leq k \leq r} \|\bar{\psi}_k^{(4)}\|_\infty \leq C_0 \quad \text{for some } 0 < C_0 < \infty.$$

SPARSE case. In this case, we only consider the situation when σ^2 is fixed (i.e., it does not vary with n). We shall first deal with the case when m_i 's are bounded. Then we extend our results to the situation when m_i 's increase slowly with sample size, and are of the same order of magnitude for all i (condition **B1**). We also assume a boundedness condition for the random design (condition **B2**).

- B1** The number of measurements m_i satisfy $\underline{m} \leq m_i \leq \bar{m}$ with $4 \leq \underline{m}$ and \bar{m}/\underline{m} is bounded by some constant $d_2 > 0$. Also, $\bar{m} = O(n^\kappa)$ for some $\kappa \geq 0$.
- B2** For each i , $\{T_{ij} : j = 1, \dots, m_i\}$ are i.i.d. from a distribution with density g , where g satisfies

$$c_{g,0} \leq g(x) \leq c_{g,1} \quad \text{for all } x \in [0, 1], \quad \text{where } 0 < c_{g,0} \leq c_{g,1} < \infty. \quad (3)$$

Finally, we have a condition on the l^2 error for approximating the covariance kernel in the model space $\mathcal{M}_{M,r}$. Define the *maximal approximation error* for an optimal parameter (B_*, Λ_*) as:

$$\bar{\beta}_n := \max_{1 \leq i \leq n} \frac{1}{m_i} \|\bar{\Sigma}_i - \Sigma_{*i}\|_F. \quad (4)$$

$$\mathbf{C} \quad \bar{m}\bar{\beta}_n = O\left(\sqrt{\frac{M \log n}{n}}\right).$$

If we use orthonormalized cubic B -spline basis for representing the eigenfunctions, then **C** follows from **A1-A2** and **B1-B2**, if the covariance kernel is indeed of rank r :

Lemma 1 : *If **A1-A2** and **B1-B2** hold, $\bar{\Sigma}_0$ is of rank r , and we use the orthonormalized cubic B -spline basis with equally spaced knots to represent the eigenfunctions, then **C** holds, if $M^{-1}(n\bar{m}^2/\log n)^{1/9} = O(1)$.*

Proof of Lemma 1 follows from the fact that for a cubic B -spline basis, for sufficiently large M , we can choose (B_*, Λ_*) such that (i) $\max_{1 \leq k \leq r} \|\bar{\psi}_k - \psi_{*k}\|_\infty = O(M^{-4})$ (by **A1** and **A2**), and (ii) $\bar{\beta}_n = O(M^{-4})$ (see Appendix A). This implies that $\|\bar{\Sigma}_0 - \Sigma_{*0}\|_F = O(M^{-4})$. The assumption that the covariance kernel is of finite rank can be relaxed somewhat by considering the true parameter

as a sequence of covariance kernels $\bar{\Sigma}_{0,n}$ such that the $(r+1)$ -th largest eigenvalue $\bar{\lambda}_{r+1,n}$ decays to zero sufficiently fast. Note that in Lemma 1, the use of B-spline basis is not essential. The result holds under the choice of any *stable basis* (i.e., the Gram matrix has a bounded condition number) with sufficient smoothness.

We now state the main result in the following theorem.

Theorem 1 (sparse case): *Suppose that **A1-A2**, **B1-B2** and **C** hold, and \bar{m} is bounded. Suppose further that M satisfies*

$$M \rightarrow \infty, \quad \text{such that} \quad M^{-1}(n/\log n)^{1/9} = O(1) \quad \text{and} \quad M = o(\sqrt{n/\log n}), \quad \text{as } n \rightarrow \infty. \quad (5)$$

Then, given $\eta > 0$, there exists $c_{0,\eta} > 0$ such that for $\alpha_n = c_{0,\eta}\sigma\sqrt{\frac{\bar{m}^2 M \log n}{n}}$, with probability at least $1 - O(n^{-\eta})$, there is a minimizer $(\hat{B}, \hat{\Lambda})$ of (2) satisfying

$$\begin{aligned} \|\hat{B} - B_*\|_F &\leq \alpha_n, \\ \|\hat{\Lambda} - \Lambda_*\|_F &\leq \alpha_n. \end{aligned}$$

Moreover, the corresponding estimate of the covariance kernel, viz., $\hat{\Sigma}_0(u, v) = \sum_{k=1}^r \hat{\lambda}_k \hat{\psi}_k(u) \hat{\psi}_k(v)$, satisfies, with probability at least $1 - O(n^{-\eta})$,

$$\|\hat{\Sigma}_0 - \bar{\Sigma}_0\|_F = O(\alpha_n).$$

Corollary 1: *Suppose that the conditions of Theorem 1 hold. Then the best rate of convergence holds if $M \asymp (n/\log n)^{1/9}$, and the corresponding rate is given by $\alpha_n \asymp (\log n/n)^{4/9}$. For estimating the eigenfunctions, this is within a factor of $\log n$ of the optimal rate. The optimal rate over a class \mathcal{C} of covariance kernels of rank r satisfying conditions **A1-A2**, and the random design points satisfying conditions **B1-B2** (with \bar{m} bounded), is $n^{-4/9}$.*

Notice that, the rate obtained here for the estimated eigenvalues is not optimal. We expect a parametric rate of convergence for the latter, which can be achieved by establishing an efficient score representation of the estimators along the line of Theorem 4. The following result generalizes Theorem 1 by allowing for m_i 's to slowly increase with n , and its proof is encapsulated in the proof of Theorem 1.

Corollary 2: *Suppose that, **A1-A2**, **B1-B2** and **C** hold. Suppose further that, \bar{m} and M satisfy*

$$\begin{aligned} (i) \quad &\bar{m}^4 M \log n = o(n), \quad (ii) \quad \max\{\bar{m}^3 M^{5/2} (\log n)^2, \bar{m}^{7/2} M^2 (\log n)^{3/2}\} = o(n), \\ (iii) \quad &M^{-1}(n\bar{m}^2/\log n)^{1/9} = O(1), \quad (iv) \quad \bar{m}^2 M^2 \log n = o(n). \end{aligned} \quad (6)$$

Then the conclusion of Theorem 1 holds. Also, the best rate is obtained when $M \asymp (n\bar{m}^2/\log n)^{1/9}$, and the corresponding $\alpha_n \asymp \bar{m}^{10/9}(\log n/n)^{4/9}$.

Condition (i) is required to ensure that $\bar{m}^2 \alpha_n^2 = o(1)$; condition (ii) is needed to ensure that the upper bound in (40) in Lemma 4 is $o(1)$; condition (iii) ensures that **C** holds; and finally, condition (iv) is used in proving Lemmas 4, 5 and 6 in Appendix B. A sufficient condition for (6) to hold is that $\bar{m} = O(n^{1/5})$ and $M \asymp (n\bar{m}^2/\log n)^{1/9}$. Notice that the best rate obtained in Corollary 2 is not optimal in general. It is near-optimal (up to a factor of $\log n$ of the optimal rate) only when \bar{m} is bounded above (Theorem 1).

DENSE case. This case refers to the scenario where the number of time points per curve is large, such that $\min_{1 \leq i \leq n} m_i \rightarrow \infty$ sufficiently fast (see condition **D** and the corresponding discussion). For simplicity, we assume further that the number of design points is the same for all the sample curves, which is not essential for the validity of the results. Denote this common value by m . In terms of the asymptotic analysis, there is an important distinction between the *sparse case* and *dense case*. For the purpose of further exposition and the proof of the result on consistency of REML estimator in the *dense case*, it is more convenient to work with the transformed data $\tilde{Y}_i = \Phi_i Y_i$. Let $\Gamma_i = \frac{1}{m} \Phi_i \Sigma_i \Phi_i^T$ and $R_i = \frac{1}{m} \Phi_i \Phi_i^T$. Then $\Gamma_i = m R_i B \Lambda B^T R_i + \sigma^2 R_i$. Then, a way of estimating $\{(\bar{\lambda}_k, \bar{\psi}_k)\}_{k=1}^r$ is by minimizing the negative log-likelihood of the transformed data:

$$\tilde{L}_n(B, \Lambda) = \frac{1}{2n} \sum_{i=1}^n \text{tr}(\Gamma_i^{-1} \frac{1}{m} \tilde{Y}_i \tilde{Y}_i^T) + \frac{1}{2n} \sum_{i=1}^n \log |\Gamma_i|. \quad (7)$$

Notice that, if R_i 's are non-singular for all i , then by direct computation, we have that the negative log-likelihoods for the raw data: (2) and that of the transformed data: (7) differ only by a constant independent of the parameters B and Λ . Hence, on the set $\{R_i \text{ are non-singular for all } i\}$, the estimators obtained by minimizing (2) and (7) are the same. Assumptions **B1** and **B2** are now replaced by:

B1' $m = O(n^\kappa)$ for some $\kappa > 0$.

D Given $\eta > 0$, there exist constants $c_{1,\eta}, c_{2,\eta} > 0$ such that

$$\mathbb{P} \left(\max_{1 \leq i \leq n} \|R_i - I_M\| \leq c_{1,\eta} \sqrt{\frac{\sigma^2}{m \log n}}, \quad \max_{1 \leq i \leq n} \|B_*^T R_i B_* - I_r\| \leq c_{2,\eta} \frac{\sigma^2}{m \log n} \right) \geq 1 - O(n^{-\eta}). \quad (8)$$

Denote the event described in (8) by $A_{1,\eta}$. Note that $A_{1,\eta}$ is defined in terms of $\mathbf{T} := \{T_{ij} : j = 1, \dots, m; i = 1, \dots, n\}$ alone. We assume throughout that $\sigma^2 \leq m$ (note that σ^2/m can be viewed as the signal-to-noise ratio). Therefore, for n large enough, on $A_{1,\eta}$, R_i is invertible for all i . The condition **D** gives concentration of individual R_i 's around the identity matrix and is discussed in more detail at the end of this section. Finally, we make an assumption about the maximal approximation error $\bar{\beta}_n$ defined through (4) which differs slightly from the condition **C** in the sparse case.

C' Given $\eta > 0$, there is a constant $c_\eta > 0$ such that $\bar{\beta}_n \leq c_\eta \frac{\sigma^2}{m} \sqrt{\frac{M \log n}{n}}$ with probability at least $1 - O(n^{-\eta})$.

A result similar to Lemma 1 can be proved to ensure condition **C'** when a stable basis is used.

Theorem 2 (dense case): Suppose that **A1-A2**, **B1'**, **C'** and **D** hold, and $m \geq \sigma^2 > 0$. Then, given $\eta > 0$, there exists $c_{0,\eta} > 0$ such that for $\alpha_n = c_{0,\eta} \sigma \sqrt{\frac{M \log n}{nm}}$, with probability at least $1 - O(n^{-\eta})$, there is a minimizer $(\hat{B}, \hat{\Lambda})$ of (7) satisfying

$$\begin{aligned} \|(I_M - B_* B_*^T)(\hat{B} - B_*)\|_F &\leq \alpha_n, \\ \|B_*^T(\hat{B} - B_*)\|_F &\leq \sqrt{\frac{m}{\sigma^2}} \alpha_n, \\ \|\hat{\Lambda} - \Lambda_*\|_F &\leq \sqrt{\frac{m}{\sigma^2}} \alpha_n. \end{aligned}$$

Further, the corresponding estimated covariance kernel $\hat{\Sigma}_0(u, v) = \sum_{k=1}^r \hat{\lambda}_k \hat{\psi}_k(u) \hat{\psi}_k(v)$ satisfies, with probability at least $1 - O(n^{-\eta})$,

$$\| \hat{\Sigma}_0 - \bar{\Sigma}_0 \|_F = O \left(\sqrt{\frac{M \log n}{n}} \right).$$

The proof of Theorem 2 requires a slight refinement of the techniques used in proving Theorem 3 stated in Section 3.2, making heavy use of condition **D**. To save space, we omit the proof. Note that the best rate in Theorem 2 implicitly depends on conditions **C'** and **D** in a complicated way, which is not the optimal rate for estimating the principal components. The optimal rate for l_2 risk of the eigenfunctions in this context is conjectured to be of the order $\max\{(\sigma^2/nm)^{8/9}, (1/n)\}$ with the second term within brackets appearing only when $r > 1$. This can be verified for the case $r = 1$ with a refinement of the proof of Corollary 1.

Discussion on condition D: We shall only consider the setting of an uniform design - either fixed, or random. The condition (8) clearly requires m to be sufficiently large, since it gives concentration of individual R_i 's around the identity matrix. To fulfil **D**, we also need some conditions on the basis functions used. Specifically, we concentrate on the following classes of basis functions. We assume that the basis functions are at least 3 times continuously differentiable.

E1 (*Sinusoidal basis*) $\max_{1 \leq k \leq M} \|\phi_k\|_\infty = O(1)$.

E2 (*Spline-type basis*) (i) For any $k \in \{1, \dots, M\}$, at most for a bounded number of basis functions ϕ_l , $\text{supp}(\phi_k) \cap \text{supp}(\phi_l)$ is nonempty; (ii) $\max_{1 \leq k \leq M} \|\phi_k\|_\infty = O(\sqrt{M})$.

One of the key observations in the case of functional data is that, the eigenfunctions $\{\psi_{*k}\}_{k=1}^r$ of the kernel Σ_{0*} (belonging to the model space) have the same degree of smoothness as the basis $\{\phi_k\}_{k=1}^M$, and the functions $\{\psi_{*k}\}_{k=1}^r$ and their derivatives are bounded. Also, notice that, $B_*^T R_i B_* = ((\frac{1}{m} \sum_{j=1}^m \psi_{*k}(T_{ij}) \psi_{*l}(T_{ij})))_{k,l=1}^r$. Based on these observations, we present some sufficient conditions for (8) to hold under the uniform design and bases of type **E1** or **E2**. We omit the proof, which uses Bernstein's inequality (in the random design case) and the *Trapezoidal rule* (in the fixed design case).

Proposition 1: Suppose that the basis is of type **E1** or **E2**. In the case of random, uniform design, (8) is satisfied if $(M \log n)^2 / \sigma^2 = O(1)$, and $\sqrt{m} \log n / \sigma^2 = O(1)$. In the case of fixed, uniform design, (8) holds (with probability 1) if $\frac{M^2 \log n}{m} (1 + \frac{M^{5/2}}{m})^2 / \sigma^2 = O(1)$, and $\log n / \sigma^2 = O(1)$. Moreover, in this setting, if the eigenfunctions $\{\psi_k\}_{k=1}^r$ vanish at the boundaries, and if the basis functions are chosen so that they also vanish at the boundaries, it is sufficient that $\frac{M^7 \log n}{m^{7/2}} / \sigma^2 = O(1)$ and $\frac{\log n}{m} / \sigma^2 = O(1)$.

Note that, two obvious implications of Proposition 1 are that (i) m needs to be rather large; and (ii) σ^2 may need to grow with n , in order that **D** holds.

Remark 1: It is to be noted that even though the consistency results for the functional data problem are proved under a specific choice of the basis for representing the eigenfunctions, viz., the (orthonormalized) cubic B-spline basis with equally spaced knots, this is by no means essential. The main features of this basis are given in terms of the various properties described in Appendix A. The crucial aspects are: (a) the basis is stable; (b) the basis functions have a certain order of

smoothness; and (c) the basis functions have fast decay away from an interval of length $O(M^{-1})$ where M is the number of basis functions used. Same consistency results can be proved as long as those properties are satisfied.

Remark 2: When \overline{m} , the number of measurements is bounded, we can relax condition **A2** to that the eigenfunctions are twice continuously differentiable and with bounded second derivative, and under this assumption we can prove a result analogous to Theorem 1 and Corollary 1, with the corresponding optimal rate of convergence being $n^{-2/5}$ instead of $n^{-4/9}$.

3 High-dimensional vector

In this section, we describe a scenario where the observations are i.i.d. Gaussian vectors, which can be approximately represented in a known lower dimensional space (see **C''**), where the effective dimensionality of the observations grows at a rate slower than the sample size. For convenience, we refer this setting as the *matrix case*. It can be seen that besides the proofs of the results derived in this section sharing a lot of common features with those in Section 2, these results also suggest an asymptotic equivalence between the *dense case* for functional data, and the *matrix case*. This means that understanding one problem helps in understanding the other problem. In particular, we conjecture that the results derived for the Gaussian vectors, such as the efficient score representation (Theorem 4), can be carried over to the functional data case with dense measurements.

3.1 Model

Suppose that we have i.i.d. observations Y_1, \dots, Y_n from $N_m(0, \overline{\Sigma})$. Assume the covariance matrix $\overline{\Sigma}$ has the following structure

$$\overline{\Sigma} = \overline{\Sigma}_0 + \sigma^2 I_m.$$

This may be regarded as a “signal-plus-noise” model, with σ^2 representing the variance of the isotropic noise component. We further assume that $\overline{\Sigma}_0$ has at least r positive eigenvalues, for some $r \geq 1$. The eigenvalues of $\overline{\Sigma}_0$ are given by $s\overline{\lambda}_1 > \dots > s\overline{\lambda}_r > s\overline{\lambda}_{r+1} \geq \dots \geq 0$, where $s > 0$ is a parameter representing the “signal strength” (so that s/σ^2 represents the signal-to-noise ratio). We assume that the observations can be well represented in a known M dimensional basis Φ with $M \leq m$ (condition **C''**). Then the model space $\mathcal{M}_{M,r}(\Phi)$ (with $r \leq M \leq m$) is defined as the set of all $m \times m$ matrices Σ of the form $\Sigma = s\Phi^T B \Lambda B^T \Phi + \sigma^2 I_m$, where Φ is an $M \times m$ matrix satisfying $\Phi\Phi^T = I_M$, $B \in \mathcal{S}_{M,r}$ and Λ is $r \times r$, diagonal with positive diagonal elements. Note that, in order to prove consistency of the REML estimator, we require that the intrinsic dimension M grows with n sufficiently slowly. In fact, it has been shown (e.g. in Paul (2005)) that, when $s/\sigma^2 = O(1)$, M must be $o(n)$ to achieve consistency.

Throughout we assume that σ^2 and s are known. Of course, we can estimate the eigenvalues of $\overline{\Sigma}$ without any knowledge of s . The unknown parameters of the model are B and Λ . The parameter space is therefore $\Omega = \mathcal{S}_{M,r} \otimes \mathbb{R}_+^r$. The estimate $(\hat{B}, \hat{\Lambda})$ of (B, Λ) is obtained by minimizing over Ω the negative log-likelihood (up to an additive constant and the multiplicative factor n),

$$L_n(B, \Lambda) = \frac{1}{2n} \text{tr}(\Sigma^{-1} \sum_{i=1}^n Y_i Y_i^T) + \frac{1}{2} \log |\Sigma|. \quad (9)$$

We then set the estimator of the first r eigenvectors of $\overline{\Sigma}$ as $\hat{\Psi} = \Phi^T \hat{B}$.

Similar to the *dense case*, for asymptotic analysis, it is more convenient to work with the transformed data $\tilde{Y}_i = \Phi Y_i$. Let $\Gamma = \Phi \Sigma \Phi^T = sB\Lambda B^T + \sigma^2 I_M$. Then one can obtain estimates of (B, Λ) by minimizing over Ω the negative log-likelihood of the transformed data:

$$\tilde{L}_n(B, \Lambda) = \frac{1}{2n} \text{tr}(\Gamma^{-1} \sum_{i=1}^n \tilde{Y}_i \tilde{Y}_i^T) + \frac{1}{2} \log |\Gamma|, \quad (10)$$

which results in the same estimate obtained by minimizing (9).

Remark 3: It is known that (Muirhead, 1982), in the setting described above, the REML estimators of (B, Λ) coincide with the first r principal components of the sample covariance matrix of $\tilde{Y}_i = \Phi Y_i$, $i = 1, \dots, n$. On the other hand, based on the calculations carried out in Appendix D, it is easy to see that the PCA estimators $(\hat{B}^{PC}, \hat{\Lambda}^{PC})$ satisfy the likelihood equations $\nabla_B \tilde{L}_n(\hat{B}^{PC}, \hat{\Lambda}^{PC}) = 0$ and $\nabla_\Lambda \tilde{L}_n(\hat{B}^{PC}, \hat{\Lambda}^{PC}) = 0$. Thus, our approach provides an independent verification of the known result that the PCA estimates are REML estimators under the rank-restricted covariance model studied here.

3.2 Consistency

We make the following assumptions about the covariance matrix.

- A1'** The eigenvalues of $\bar{\Sigma}_0$ are given by $s\bar{\lambda}_1 \geq \dots \geq s\bar{\lambda}_m \geq 0$ and satisfy, for some $r \geq 1$ (fixed),
(i) $\bar{c}_1 \geq \bar{\lambda}_1 > \dots > \bar{\lambda}_r > \bar{\lambda}_{r+1}$ for some $\bar{c}_1 < \infty$; (ii) $\max_{1 \leq j \leq r} (\bar{\lambda}_j - \bar{\lambda}_{j+1})^{-1} \leq \bar{c}_2 < \infty$.
- C''** Assume that there exists $(B_*, \Lambda_*) \in \Omega$ (referred as “optimal parameter”) such that, the matrix $\Sigma_{*0} = s\Phi^T B_* \Lambda_* B_*^T \Phi$ is a close approximation to $\bar{\Sigma}_0$ in the sense that $\beta_n := \|\bar{\Sigma}_0 - \Sigma_{*0}\|_F = O(\sigma^2 \sqrt{\frac{M \log n}{n}})$.

Note that **C''** implies that the observation vectors can be closely approximated in the basis Φ .

Theorem 3 (matrix case): Suppose that **A1'** and **C''** hold, and $s \geq \sigma^2 > 0$. Then given $\eta > 0$, there exists $c_{0,\eta} > 0$ such that for $\alpha_n = c_{0,\eta} \sigma \sqrt{\frac{M \log n}{ns}}$, with probability at least $1 - O(n^{-\eta})$, there is a minimizer $(\hat{B}, \hat{\Lambda})$ of (10) satisfying

$$\begin{aligned} \|(I_M - B_* B_*^T)(\hat{B} - B_*)\|_F &\leq \alpha_n, \\ \|B_*^T(\hat{B} - B_*)\|_F &\leq \sqrt{\frac{s}{\sigma^2}} \alpha_n, \\ \|\hat{\Lambda} - \Lambda_*\|_F &\leq \sqrt{\frac{s}{\sigma^2}} \alpha_n. \end{aligned}$$

Observe that the rates obtained in Theorem 2 and Theorem 3 are identical once we replace m in Theorem 2 by s . Thus the number of measurements m in the *dense case* is an analog of the signal strength s in the *matrix case*. This important observation suggests an *asymptotic equivalence* between these two problems. This is a result of the concentration of the matrices $\{R_i\}_{i=1}^n$ around I_M for the *dense case* (condition **D**). Under the matrix case, the analogs of R_i exactly equal the identity matrix. Moreover, Theorem 3 establishes the closeness of the REML estimator to the optimal parameter, which serves as an important step towards proving Theorem 4.

3.3 Efficient score representation

When the observations are i.i.d. Gaussian vectors, we can get a more refined result than the one stated in Theorem 3. In this section, we show that by using the intrinsic geometry, we can get an *efficient score representation* of the REML estimator (and hence PCA estimator). In Paul and Johnstone (2007), a first order approximation to the sample eigenvectors (i.e. PCA estimates) is obtained using matrix perturbation theory (Kato, 1980). Subsequently, it has also been shown there that the rate of convergence of l^2 -risk of PCA estimators is optimal. Here, we show that the efficient score representation of the REML estimator coincides with this first order approximation when the signal-to-noise ratio s/σ^2 is bounded (Corollary 3). Our approach is different from the perturbation analysis. It also quantifies the role of intrinsic geometry of the parameter space explicitly. Our result gives an alternative interpretation of this approximation, and consequently, the score representation points to an asymptotic optimality of the REML (and hence PCA) estimator.

We first introduce some notations. More details can be found in Appendix D. Let $\zeta = \log \Lambda$ (treated interchangeably as an $r \times 1$ vector and an $r \times r$ diagonal matrix). The parameter space for (B, ζ) is $\tilde{\Omega} := \mathcal{S}_{M,r} \otimes \mathbb{R}^r$. Let $\mathcal{T}_B := \{U \in \mathbb{R}^{M \times r} : B^T U = -U^T B\}$ denote the tangent space of the *Stiefel manifold* $\mathcal{S}_{M,r}$ at B . Then the tangent space for the product manifold $\tilde{\Omega}$ at (B, ζ) is $\mathcal{T}_B \oplus \mathbb{R}^r$ (see Appendix E for the definition of the product manifold and its tangent space).

For notational simplicity, we use θ_* to denote (B_*, ζ_*) and θ_0 to denote (B_0, ζ_0) . Define $L(\theta_0; \theta_*) = \mathbb{E}_{\theta_*} \tilde{L}_n(\theta_0)$. Let $\nabla \tilde{L}_n(\cdot)$ and $\nabla L(\cdot; \theta_*)$ denote the *intrinsic gradient* of the functions $\tilde{L}_n(\cdot)$ and $L(\cdot; \theta_*)$ with respect to (B, ζ) , respectively. Also, let $H_n(\cdot)$ and $H(\cdot; \theta_*)$ denote the *intrinsic Hessian operator* of the functions $\tilde{L}_n(\cdot)$ and $L(\cdot; \theta_*)$ with respect to (B, ζ) , respectively. Let $H^{-1}(\cdot; \theta_*)$ denote the inverse Hessian operator of $L(\cdot; \theta_*)$. Also we use $H_B(\cdot; \theta_*)$ to denote the Hessian of $L(\cdot; \theta_*)$ w.r.t. B . Notations for Hessian w.r.t. ζ and gradients w.r.t B and ζ are defined similarly.

The following result gives the efficient score representation of the REML estimator in the situation when $\sigma^2 = 1$ and $s = 1$. The result can be extended via rescaling to the case for arbitrary σ^2 and s with $s \geq \sigma^2 > 0$, and s/σ^2 being bounded.

Theorem 4 (score representation): *Suppose that **A1'** and **C''** hold with $\sigma^2 = 1$, $s = 1$, and $M = o(n^a)$ for some $a \in (0, 1)$. Let $\gamma_n = \max\{\sqrt{\frac{M \vee \log n}{n}}, \beta_n\}$. Then there is a minimizer $(\hat{B}, \hat{\Lambda})$ of the negative log-likelihood (10) such that, with probability tending towards 1,*

$$\hat{B} - B_* = -H_B^{-1}(\theta_*; \theta_*)(\nabla_B \tilde{L}_n(\theta_*)) + O(\gamma_n^2) \quad (11)$$

$$\hat{\Lambda} - \Lambda_* = -\Lambda_* H_\zeta^{-1}(\theta_*; \theta_*)(\nabla_\zeta \tilde{L}_n(\theta_*)) + O(\gamma_n^2). \quad (12)$$

In particular, from this representation, we have, with probability tending towards 1,

$$\|\hat{B} - B_*\|_F = O(\gamma_n); \quad (13)$$

$$\|\hat{\Lambda} - \Lambda_*\|_F = O(\sqrt{\frac{\log n}{n}} + \gamma_n^2). \quad (14)$$

Note that Theorem 4 gives the optimal rate of convergence for l^2 -risk of the estimated eigenvectors when $\beta_n = 0$ (i.e., no model bias). This result follows from the minimax lower bound on the risk obtained by Paul and Johnstone (2007). Note that, this lower bound under the current setting follows essentially from the proof of Corollary 1. Also, when $a \leq 1/2$, this result shows that $\hat{\Lambda}$ converges at a parametric rate. Indeed, the representation (12) implies asymptotic normality of $\hat{\Lambda}$

when $a \leq 1/2$. In the derivation of Theorem 4 we need to compute the Hessian and its inverse, which leads to the following representation.

Corollary 3: *Under the assumptions of Theorem 4, we have the following representation:*

$$H_B^{-1}(\theta_*; \theta_*)(\nabla_B \tilde{L}_n(\theta_*)) = [\mathbf{R}_1 \tilde{S} B_{*1} : \cdots : \mathbf{R}_r \tilde{S} B_{*r}],$$

where B_{*j} is the j -th column of B_* , and

$$\mathbf{R}_j = \sum_{1 \leq i \neq j \leq r} \frac{1}{(\lambda_{*i} - \lambda_{*j})} B_{*i} B_{*i}^T - \frac{1}{\lambda_{*j}} (I_M - B_* B_*^T),$$

is the resolvent operator corresponding to Γ_* “evaluated at” $(1 + \lambda_{*j})$.

Combining Corollary 3 with (11), we get a first order approximation to \hat{B} which coincides with the approximation for sample eigenvectors obtained in Paul and Johnstone (2007). However, Theorem 4 has deeper implications. Since it gives an efficient score representation, it suggests an asymptotic optimality of the REML estimators in the minimax sense.

4 Proof of Theorem 1

Since σ^2 is fixed and assumed known, without loss of generality, we take $\sigma^2 = 1$. In this section, we give an outline of the main ideas/steps. The details of the proofs are given in Appendix B. The strategy of the proof is as follows. We restrict our attention to a subset $\Theta(\alpha_n)$ of the parameter space (referred as the *restricted parameter space*), which is the image under exponential map of the boundary of an ellipsoid centered at 0, in the tangent space of an “optimal parameter”. We then show that with probability tending towards 1, for every parameter value in this restricted parameter space, the value of the negative log-likelihood is greater than the value of the negative log-likelihood at the optimal parameter. Due to the Euclidean geometry of the tangent space, this implies that with probability tending towards 1, there is a local maximum of the log-likelihood within the image (under exponential map) of the closed ellipsoid. The key steps of the proof are:

- (i) Decompose the difference between the negative log-likelihood at the optimal parameter and an arbitrary parameter in the restricted space as a sum of three terms - a term representing the average Kullback-Leibler divergence between the distributions, a term representing random fluctuation in the log-likelihood, and a term representing the *model bias* (equation (19)).
- (ii) For every fixed parameter in the restricted parameter space: (a) provide upper and lower bounds (dependent on α_n) for the average Kullback-Leibler divergence; (b) provide upper bounds for the random term and the model bias term. In both cases, the bounds are probabilistic with exponentially small tails.
- (iii) Use a covering argument combined with a union bound to extend the above probabilistic bounds on difference between log-likelihoods corresponding to a single parameter in $\Theta(\alpha_n)$ to the infimum of the difference over the entire $\Theta(\alpha_n)$.

The strategy of this proof is standard. However, in order to carry it out we need to perform detailed computations involving the geometry of the parameter space such as the structure of the tangent space and the exponential map. Note that, in the current case the geometry of the parameter

space is well-understood, so that there exist explicit form of the exponential map and a precise description of the tangent space. This helps in obtaining the precise form of the local Euclidean approximations around an optimal parameter in the derivations.

4.1 Parameter space and exponential map

We use the following characterization of the tangent space \mathcal{T}_B of the Stiefel manifold $\mathcal{S}_{M,r}$ at a point B . Any element $U \in \mathcal{T}_B$ can be expressed as $U = BA_U + C_U$, where $A_U = -A_U^T$ and $B^T C_U = O$. We then define the restricted parameter space centered at an optimal parameter (B_*, Λ_*) by

$$\begin{aligned} \Theta(\alpha_n) &:= \{(\mathbf{exp}(1, B_* A_U + C_U), \Lambda_* \exp(D)) : A_U = -A_U^T, B_*^T C_U = O, D \in \mathbb{R}^r, \\ &\quad \text{such that } \|A_U\|_F^2 + \|C_U\|_F^2 + \|D\|_F^2 \leq \alpha_n^2\}. \end{aligned} \quad (15)$$

In the definition of $\Theta(\alpha_n)$ and henceforth, we shall treat Λ and D interchangeably as an $r \times 1$ vector, and an $r \times r$ diagonal matrix. The function $\mathbf{exp}(t, U)$ is the *exponential map* on $\mathcal{S}_{M,r}$ at B_* , mapping a tangent vector in \mathcal{T}_{B_*} to a point on the manifold. For $U \in \mathcal{T}_{B_*}$ and $t \geq 0$, it is defined as

$$\mathbf{exp}(t, U) = B_* \mathbf{M}(t, U) + Q \mathbf{N}(t, U), \quad \text{where} \quad \begin{bmatrix} \mathbf{M}(t, U) \\ \mathbf{N}(t, U) \end{bmatrix} = \exp \left(t \begin{bmatrix} B_*^T U & -R^T \\ R & O \end{bmatrix} \right) \begin{bmatrix} I_r \\ O \end{bmatrix},$$

where $\exp(\cdot)$ is the usual matrix exponential, and $QR = (I_M - B_* B_*^T)U$ is the QR-decomposition. The properties of the map $\mathbf{exp}(1, \cdot)$ that we shall heavily use in the subsequent analysis (see Appendix B) are : for $U \in \mathcal{T}_{B_*}$,

$$B_*^T (\mathbf{exp}(1, U) - B_*) = B_*^T U + O(\|B_*^T U\|_F + \|(I_M - B_* B_*^T)U\|_F) \|U\|_F, \quad (16)$$

$$(I_M - B_* B_*^T) \mathbf{exp}(1, U) = (I_M - B_* B_*^T)U + O(\|(I_M - B_* B_*^T)U\|_F) \|U\|_F, \quad (17)$$

as $\|U\|_F \rightarrow 0$. These properties are easily verified by using the definition of the matrix exponential $\exp(\cdot)$, and the Taylor series expansion.

4.2 Loss decomposition

We shall show that, given $\eta > 0$, for an appropriate choice of the constant $c_{0,\eta}$ in the definition of α_n (in Theorem 1), for large enough n , we have

$$\mathbb{P} \left(\inf_{(B, \Lambda) \in \Theta(\alpha_n)} L_n(B, \Lambda) > L_n(B_*, \Lambda_*) \right) \geq 1 - O(n^{-\eta}). \quad (18)$$

From this, it follows immediately that with probability tending towards 1, there is a local minimum $(\hat{B}, \hat{\Lambda})$ of $L_n(B, \Lambda)$ in the set $\bar{\Theta}(\alpha_n)$ defined as

$$\begin{aligned} \bar{\Theta}(\alpha_n) &= \{(\mathbf{exp}(1, B_* A_U + C_U), \Lambda_* \exp(D)) : A_U = -A_U^T, B_*^T C_U = O, D \in \mathbb{R}^r \\ &\quad \text{such that } \|A_U\|_F^2 + \|C_U\|_F^2 + \|D\|_F^2 \leq \alpha_n^2\}, \end{aligned}$$

which concludes the proof of Theorem 1.

We start with the basic decomposition:

$$\begin{aligned} &L_n(B, \Lambda) - L_n(B_*, \Lambda_*) \\ &= [\mathbb{E}L_n(B, \Lambda) - \mathbb{E}L_n(B_*, \Lambda_*)] + [(L_n(B, \Lambda) - \mathbb{E}L_n(B, \Lambda)) - (L_n(B_*, \Lambda_*) - \mathbb{E}L_n(B_*, \Lambda_*))] \\ &= \frac{1}{n} \sum_{i=1}^n K(\Sigma_i, \Sigma_{*i}) + \frac{1}{2n} \sum_{i=1}^n \text{tr}((\Sigma_i^{-1} - \Sigma_{*i}^{-1})(S_i - \bar{\Sigma}_i)) + \frac{1}{2n} \sum_{i=1}^n \text{tr}((\Sigma_i^{-1} - \Sigma_{*i}^{-1})(\bar{\Sigma}_i - \Sigma_{*i})), \end{aligned} \quad (19)$$

where $S_i = Y_i Y_i^T$ and

$$K(\Sigma_i, \Sigma_{*i}) = \frac{1}{2} \text{tr} \left(\Sigma_i^{-1/2} (\Sigma_{*i} - \Sigma_i) \Sigma_i^{-1/2} \right) - \frac{1}{2} \log |I_{m_i} + \Sigma_i^{-1/2} (\Sigma_{*i} - \Sigma_i) \Sigma_i^{-1/2}|,$$

is the Kullback-Leibler divergence corresponding to observation i . Note that the proofs of Theorems 2 and 3 share a lot of commonality with the *sparse case* discussed here, in that these proofs depend on the same basic decomposition of the loss function.

4.3 Probabilistic bounds for a fixed parameter in $\Theta(\alpha_n)$

In order to derive the results in the following three propositions, we need to restrict our attention to an appropriate subset of the space of the design points \mathbf{T} which has high probability. Accordingly, given $\eta > 0$, we define such a set A_η through (34) in Proposition 8 (in Appendix A). The following proposition gives probabilistic bounds for the average Kullback-Leibler divergence in terms of α_n .

Proposition 2: *Given $\eta > 0$, for every $(B, \Lambda) \in \Theta(\alpha_n)$, there is a set $A_{1,\eta}^{B,\Lambda}$ (depending on (B, Λ)), defined as*

$$A_{1,\eta}^{B,\Lambda} := \left\{ d'_\eta \alpha_n^2 \leq \frac{1}{n} \sum_{i=1} K(\Sigma_i, \Sigma_{*i}) \leq d''_\eta \bar{m}^2 \alpha_n^2 \right\}, \quad (20)$$

for appropriate positive constants d'_η and d''_η (depending on $\bar{\lambda}_1$ and r), such that for n large enough, $\mathbb{P}(A_\eta \cap (A_{1,\eta}^{B,\Lambda})^c) = O(n^{-(2+2\kappa)Mr-\eta})$.

Note that the bound in (20) is not sharp when $\bar{m} \rightarrow \infty$, which leads the suboptimal rates in Corollary 2. The following propositions bound the random term and the bias term in (19), respectively.

Proposition 3: *Given $\eta > 0$, for each $(B, \Lambda) \in \Theta(\alpha_n)$, there is a set $A_{2,\eta}^{B,\Lambda}$, defined as,*

$$A_{2,\eta}^{B,\Lambda} = \left\{ \left| \frac{1}{2n} \sum_{i=1}^n \text{tr}((\Sigma_i^{-1} - \Sigma_{*i}^{-1})(S_i - \bar{\Sigma}_i)) \right| \leq d_\eta \bar{m} \alpha_n \sqrt{\frac{M \log n}{n}} \right\},$$

for some $d_\eta > 0$, such that, $\mathbb{P}(A_{1,\eta} \cap (A_{2,\eta}^{B,\Lambda})^c) = O(n^{-(2+2\kappa)Mr-\eta})$.

Proposition 4: *Given $\eta > 0$, for each $(B, \Lambda) \in \Theta(\alpha_n)$, there is a set $A_{3,\eta}^{B,\Lambda}$, defined as,*

$$A_{3,\eta}^{B,\Lambda} = \left\{ \left| \frac{1}{2n} \sum_{i=1}^n \text{tr}[(\Sigma_i^{-1} - \Sigma_{*i}^{-1})(\bar{\Sigma}_i - \Sigma_{*i})] \right| \leq d_\eta \bar{m} \alpha_n \sqrt{\frac{M \log n}{n}} \right\},$$

for some constant $d_\eta > 0$, such that for large enough n , $\mathbb{P}(A_\eta \cap (A_{3,\eta}^{B,\Lambda})^c) = O(n^{-(2+2\kappa)Mr-\eta})$.

Combining Propositions 2-4, we obtain that, given $\eta > 0$, there is a constant $c_{0,\eta}$, such that, for every $(B, \Lambda) \in \Theta(\alpha_n)$,

$$\mathbb{P}(\{L_n(B, \Lambda) - L_n(B_*, \Lambda_*) \leq \frac{1}{2} \alpha_n^2\} \cap A_\eta) = O(n^{-(2+2\kappa)Mr-\eta}). \quad (21)$$

4.4 Covering of the space $\Theta(\alpha_n)$

To complete the proof of Theorem 1, we construct a δ_n -net in the set $\Theta(\alpha_n)$, for some $\delta_n > 0$ sufficiently small. This means that, for any $(B_1, \Lambda_1) \in \Theta(\alpha_n)$ there exists an element (B_2, Λ_2) of the net (with $B_k = \mathbf{exp}(1, B_* A_{U_k} + C_{U_k})$ and $\Lambda_k = \Lambda_* \exp(D_k)$, $k = 1, 2$), such that we have $\|B_1 - B_2\|_F^2 + \|\Lambda_1 - \Lambda_2\|_F^2 \leq \delta_n^2$. The spaces $\{A \in \mathbb{R}^{r \times r} : A = -A^T\}$ and $\{C \in \mathbb{R}^{M \times r} : B_*^T C = O\}$ are Euclidean subspaces of dimension $r(r-1)/2$ and $Mr - r^2$, respectively. Therefore, $\Theta(\alpha_n)$ is the image under $(\mathbf{exp}(1, \cdot), \exp(\cdot))$ of a hyper-ellipse of dimension $p = Mr - r(r+1)/2$. Thus, using standard construction of nets on spheres in \mathbb{R}^p , we can find such a δ_n -net $\mathcal{C}[\delta_n]$, with at most $d_1 \max\{1, (\alpha_n \delta_n^{-1})^p\}$ elements, for some $d_1 < \infty$.

If we take $\delta_n = (\overline{m}^2 n)^{-1}$, then from (21) using union bound it follows that, for n large enough,

$$\mathbb{P}\left(\left\{\inf_{(B, \Lambda) \in \mathcal{C}[\delta_n]} L_n(B, \Lambda) - L_n(B_*, \Lambda_*) > \frac{1}{2} \alpha_n^2\right\} \cap A_\eta\right) \geq 1 - O(n^{-\eta}).$$

This result, together with the following lemma and the fact that $\mathbb{P}(A_\eta) \geq 1 - O(n^{-\eta})$ (Proposition 8), as well as the definition of $\mathcal{C}[\delta_n]$, proves (18). The proof of Lemma 2 is given in Appendix B.

Lemma 2: *Let (B_k, Λ_k) , $k = 1, 2$, be any two elements of $\Theta(\alpha_n)$ satisfying $\|B_1 - B_2\|_F^2 + \|\Lambda_1 - \Lambda_2\|_F^2 \leq \delta_n^2$, with $\delta_n = (\overline{m}^2 n)^{-1}$. Then, given $\eta > 0$, there are constants $d_{3,\eta}, d_{4,\eta} > 0$, such that, the set $A_{4,\eta} := \left\{\max_{1 \leq i \leq n} \|\overline{\Sigma}_i^{-1/2} S_i \overline{\Sigma}_i^{-1/2} - I_{m_i}\|_F \leq d_{3,\eta} \overline{m} \log n\right\}$ satisfies $\mathbb{P}(A_{4,\eta} | \mathbf{T}) \geq 1 - O(n^{-\eta-1})$, for $\mathbf{T} \in A_\eta$; and on $A_{4,\eta}$, we have $|L_n(B_1, \Lambda_1) - L_n(B_2, \Lambda_2)| = o(\alpha_n^2)$.*

5 Proof of Theorem 3

There is essentially only one step where the proof of Theorem 3 differs from that of Theorem 1. It involves providing sharper bounds for the Kullback-Leibler divergence between an “optimal parameter”, and an arbitrary parameter in the restricted parameter space $\tilde{\Theta}(\alpha_n)$, an ellipsoid in the tangent space at the “optimal parameter”:

$$\begin{aligned} \tilde{\Theta}(\alpha_n) &= \{(\mathbf{exp}(1, B_* A_U + C_U), \Lambda_* \exp(D)) : A_U = -A_U^T, B_*^T C_U = O, D \in \mathbb{R}^r \\ &\quad \text{such that } \frac{\sigma^2}{s} \|A_U\|_F^2 + \|C_U\|_F^2 + \frac{\sigma^2}{s} \|D\|_F^2 = \alpha_n^2\}. \end{aligned}$$

Note that now the restricted parameter space is the image (under exponential maps) of an ellipse, whose principal axes can differ substantially depending on the signal-to-noise ratio s/σ^2 . This is crucial for obtaining the sharper bounds for the Kullback-Leibler divergence (see equation (58)). As in Section 4, our strategy is to show that, given $\eta > 0$, for an appropriate choice of $c_{0,\eta}$, for large enough n , we have

$$\mathbb{P}\left(\inf_{(B, \Lambda) \in \tilde{\Theta}(\alpha_n)} \tilde{L}_n(B, \Lambda) > \tilde{L}_n(B_*, \Lambda_*)\right) \geq 1 - O(n^{-\eta}).$$

From this, we conclude the proof of Theorem 3 using similar arguments as in the proof of Theorem 1.

Define $\tilde{S} = \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \tilde{Y}_i^T$, where $\tilde{Y}_i = \Phi Y_i$. Then, for an arbitrary $(B, \Lambda) \in \tilde{\Theta}(\alpha_n)$, we have the following decomposition:

$$\begin{aligned} \tilde{L}_n(B, \Lambda) - \tilde{L}_n(B_*, \Lambda_*) &= \left[\left(\tilde{L}_n(B, \Lambda) - \mathbb{E} \tilde{L}_n(B, \Lambda) \right) - \left(\tilde{L}_n(B_*, \Lambda_*) - \mathbb{E} \tilde{L}_n(B_*, \Lambda_*) \right) \right] \\ &= K(\Gamma, \Gamma_*) + \frac{1}{2} \text{tr} \left((\Gamma^{-1} - \Gamma_*^{-1})(\tilde{S} - \bar{\Gamma}) \right) + \frac{1}{2} \text{tr} \left((\Gamma^{-1} - \Gamma_*^{-1})(\bar{\Gamma} - \Gamma_*) \right) \end{aligned} \quad (22)$$

with

$$\begin{aligned} K(\Gamma, \Gamma_*) &= \frac{1}{2} \text{tr} (\Gamma^{-1}(\Gamma_* - \Gamma)) - \frac{1}{2} \log |I_M + \Gamma^{-1}(\Gamma_* - \Gamma)| \\ &= \frac{1}{2} \text{tr} (\Gamma^{-1/2}(\Gamma_* - \Gamma)\Gamma^{-1/2}) - \frac{1}{2} \log |I_M + \Gamma^{-1/2}(\Gamma_* - \Gamma)\Gamma^{-1/2}|, \end{aligned}$$

being the Kullback-Leibler divergence between the probability distributions $N_M(0, \Gamma)$ and $N_M(0, \Gamma_*)$, where $\Gamma^{-1/2} = (\Gamma^{1/2})^{-1}$, and $\Gamma^{1/2}$ is a symmetric, positive definite, square root of Γ . The following is an analogue of Proposition 2.

Proposition 5: *Under the assumptions of Theorem 3, there exist constants $c', c'' > 0$ such that, for sufficiently large n ,*

$$c' \alpha_n^2 \left(\frac{s}{\sigma^2} \right) \leq K(\Gamma, \Gamma_*) \leq c'' \alpha_n^2 \left(\frac{s}{\sigma^2} \right), \quad (23)$$

for all $(B, \Lambda) \in \tilde{\Theta}(\alpha_n)$, where $\Gamma = sB\Lambda B^T + \sigma^2 I_M$ and $\Gamma_* = sB_*\Lambda_* B_*^T + \sigma^2 I_M$.

The following are analogues of the Propositions 3 and 4, respectively.

Proposition 6: *Given $\eta > 0$, there exists a constant $c_\eta > 0$, such that for each $(B, \Lambda) \in \tilde{\Theta}(\alpha_n)$,*

$$\mathbb{P} \left(\left| \text{tr}((\Gamma^{-1} - \Gamma_*^{-1})(\tilde{S} - \bar{\Gamma})) \right| \leq c_\eta \sqrt{\frac{M \log n}{n}} \sqrt{\frac{s}{\sigma^2} \alpha_n} \right) \geq 1 - O(n^{-(2+2\kappa)Mr-\eta}).$$

This proposition can be easily proved using an exponential inequality by Davidson and Szarek (2001) on the fluctuations of the extreme eigenvalues of a Wishart matrix.

Proposition 7: *There is a constant $c > 0$ such that, uniformly over $(B, \Lambda) \in \tilde{\Theta}(\alpha_n)$,*

$$\left| \text{tr}((\Gamma^{-1} - \Gamma_*^{-1})(\bar{\Gamma} - \Gamma_*)) \right| \leq \| \Gamma^{-1} - \Gamma_*^{-1} \|_F \| \bar{\Gamma} - \Gamma_* \|_F \leq c \frac{1}{\sigma^2} \sqrt{\frac{s}{\sigma^2} \alpha_n \beta_n}.$$

Propositions 5-7 (together with conditions **A1'** and **C''**) show that, for an appropriate choice of $c_{0,\eta}$, $\tilde{L}_n(B, \Lambda) - \tilde{L}_n(B_*, \Lambda_*) \geq c' \alpha_n^2$, for some $c' > 0$ with very high probability, for every fixed $(B, \Lambda) \in \tilde{\Theta}(\alpha_n)$. The proof of Theorem 3 is finished by constructing a δ_n -net similarly as in Section 4.4 for the *sparse case*.

6 Proof of Theorem 4

The basic strategy of the proof is similar to that in classical inference with Euclidean parameter space. The main difficulty in present context lies in dealing with the Hessian operator of the log-likelihood (intrinsic Fisher information operator) and its inverse. Details of these calculations are given in Appendix D.

Rewrite the negative log-likelihood (10) (up to a multiplicative constant) as

$$\tilde{L}_n(B, \Lambda) = \text{tr}(\Gamma^{-1}\tilde{S}) + \log |\Gamma|, \quad \text{where} \quad \tilde{S} = \frac{1}{n} \sum_{i=1}^n \tilde{Y}_i \tilde{Y}_i^T. \quad (24)$$

By Theorem 3, given $\eta > 0$, there is a constant $c_{3,\eta} > 0$ such that the set

$$\tilde{A}_{3,\eta} := \left\{ \|\hat{U}\|_F^2 + \|\hat{D}\|_F^2 \leq c_{3,\eta} \alpha_n^2 \right\}$$

has probability at least $1 - O(n^{-\eta})$, where $\alpha_n = c_{0,\eta} \sqrt{\frac{M \log n}{n}}$, and $(\hat{U}, \hat{D}) \in \mathcal{T}_{B_*} \oplus \mathbb{R}^r$ is such that $(\hat{B}, \hat{\Lambda}) := (\mathbf{exp}(1, \hat{U}), \Lambda_* \mathbf{exp}(\hat{D}))$ is a minimizer of (24).

First, by the same concentration bound for singular values of random matrices with i.i.d. Gaussian entries (Davidson and Szarek, 2001) used in the proof of Proposition 6, there exists $c_{4,\eta} > 0$, such that the set

$$\tilde{A}_{4,\eta} := \left\{ \|\tilde{S} - \bar{\Gamma}\| \leq c_{4,\eta} \sqrt{\frac{M \vee \log n}{n}} \right\}$$

has probability at least $1 - O(n^{-\eta})$. It then follows that, we can choose an appropriate constant $c_{5,\eta} > 0$ such that, on $\tilde{A}_{3,\eta} \cap \tilde{A}_{4,\eta}$, $\|\nabla \tilde{L}_n(\theta_*)\| \leq c_{5,\eta} \gamma_n$, where $\gamma_n = \max\{\sqrt{\frac{M \vee \log n}{n}}, \beta_n\}$ and $\theta_* = (B_*, \Lambda_*)$. Next, for any $X = (X_B, X_\zeta) \in \mathcal{T}_{B_*} \oplus \mathbb{R}^r$, define

$$\|X\| := \left[\|X_B\|_F^2 + \|X_\zeta\|_F^2 \right]^{1/2}.$$

Also, let $\langle \cdot, \cdot \rangle_g$ denote the canonical metric on $\mathcal{T}_{B_*} \oplus \mathbb{R}^r$ (see Appendix D). Using the fact that $\nabla \tilde{L}_n(\hat{\theta}) = 0$, where $\hat{\theta} = (\hat{B}, \hat{\Lambda})$, and defining $\hat{\Delta} := (\hat{U}, \hat{D})$, then on $\tilde{A}_{3,\eta} \cap \tilde{A}_{4,\eta}$, for any $X \in \mathcal{T}_{B_*} \oplus \mathbb{R}^r$ with $\|X\| \leq 1$,

$$\begin{aligned} -\langle \nabla \tilde{L}_n(\theta_*), X \rangle_g &= \langle \nabla \tilde{L}_n(\hat{\theta}) - \nabla \tilde{L}_n(\theta_*), X \rangle_g \\ &= \langle H_n(\theta_*)(\hat{\Delta}), X \rangle_g + O(\|\hat{\Delta}\|^2) + O(\gamma_n \|\hat{\Delta}\|) \\ &= \langle H(\theta_*; \theta_*)(\hat{\Delta}), X \rangle_g + \langle [H_n(\theta_*) - H(\theta_*; \theta_*)](\hat{\Delta}), X \rangle_g + O(\alpha_n^2 + \alpha_n \gamma_n) \end{aligned} \quad (25)$$

where $H_n(\cdot)(\hat{\Delta})$ and $H(\cdot; \theta_*)(\hat{\Delta})$ are the corresponding *covariant derivatives* of $\tilde{L}_n(\cdot)$ and $L(\cdot; \theta_*)$ in the direction of $\hat{\Delta}$. By simple calculations based on the expressions in Appendix D, there exists a constant $c_{6,\eta} > 0$, such that on $\tilde{A}_{3,\eta} \cap \tilde{A}_{4,\eta}$, $\|H_n(\theta_*)(\hat{\Delta}) - H(\theta_*; \theta_*)(\hat{\Delta})\| \leq c_{6,\eta} \alpha_n \gamma_n$. It can be checked using assumptions **A1'** and **C''** that the linear operator $H^{-1}(\theta_*; \theta_*) : \mathcal{T}_{B_*} \oplus \mathbb{R}^r \rightarrow \mathcal{T}_{B_*} \oplus \mathbb{R}^r$, is bounded in operator norm (Appendix D). Therefore, using the definition of covariant derivative and inverse of Hessian, from (25) we have, on $\tilde{A}_{3,\eta} \cap \tilde{A}_{4,\eta}$,

$$\hat{\Delta} = -H^{-1}(\theta_*; \theta_*)(\nabla \tilde{L}_n(\theta_*)) + O(\alpha_n \gamma_n) + O(\alpha_n^2). \quad (26)$$

Hence, on $\tilde{A}_{3,\eta} \cap \tilde{A}_{4,\eta}$, the bound on $\|\hat{\Delta}\|$ can be improved from $O(\alpha_n)$ to

$$\|\hat{\Delta}\| = O(\gamma_n \alpha_n + \alpha_n^2). \quad (27)$$

We can then repeat exactly the same argument, by using (25) to derive (26), but now with the bound on $\|\hat{\Delta}\|$ given by (27). Since $M = O(n^a)$ for some $a < 1$, so that $\alpha_n^2 = o(\gamma_n)$, this way we get the more precise expression,

$$\hat{\Delta} = -H^{-1}(\theta_*; \theta_*)(\nabla \tilde{L}_n(\theta_*)) + O(\gamma_n^2). \quad (28)$$

Moreover, it can be easily verified that,

$$\frac{\partial}{\partial \zeta} \nabla_B L(\theta_*; \theta_*) := \mathbb{E}_{\theta_*} \left[\frac{\partial}{\partial \zeta} \nabla_B \tilde{L}_n(\theta_*) \right] = 0.$$

Hence, by (70) in Appendix E, the Hessian operator, and its inverse, are “block diagonal”, on the parameter space (viewed as a product manifold), with diagonal blocks corresponding to Hessians (inverse Hessians) w.r.t. B and ζ , respectively. This yields (11) and (12) in Theorem 4. Also, (13) and (14) follow immediately from (28).

7 Discussion

In this paper, we have demonstrated the effectiveness of utilizing the geometry of the non-Euclidean parameter space in determining consistency and rates of convergence of the REML estimators of principal components. We first study the REML estimators of eigenvalues and eigenfunctions of the covariance kernel for functional data, estimated from sparse, irregular measurements. The convergence rate of the estimated eigenfunctions is shown to be near-optimal when the number of measurements per curve is bounded and when M , the number of basis functions, varies with n at an appropriate rate (Theorem 1 and Corollary 1). The technique used in proving Theorem 1 is most suitable for dealing with the very sparse case (i.e., number of measurements per curve is bounded). We have also used it to prove consistency for the case where the number of measurements increases slowly with sample size (Corollary 2). However, this does not result in the optimal convergence rate. The latter case is more difficult because of the complications of dealing with inverses of random matrices (Σ_i) of growing dimensions. A more delicate analysis, that can handle this issue more efficiently, is likely to give tighter bounds for the average Kullback-Leibler divergence than that obtained in Proposition 2. Then it may be possible to extend the current technique to prove optimality of the REML estimators in a broader regime. A variant of the technique used for proving Theorem 1 also gives consistency of the REML estimator for functional data in a regime of dense measurements, as well as for a class of high-dimensional Gaussian vectors (Theorems 2 and 3). In the latter case, we also derive an efficient score representation (Theorem 4), which involves determining the intrinsic Fisher information operator and its inverse.

Now we present some conjectures we aim to pursue. First, as discussed earlier, based on the score representation, we conjecture the asymptotic optimality of the REML estimator for the *matrix case*. Secondly, we conjecture that there exists an efficient score representation of the REML estimator in the functional data problem as well. If so, then this estimator is likely to achieve the optimal nonparametric rate (for a broader regime), and may even be asymptotically optimal. This may explain the superior numerical performance of the REML estimator observed by Peng and Paul (2007). Thirdly, our results (Theorems 2 and 3) give a strong indication of an asymptotic equivalence between two classes of problems : statistical inference for functional data with dense measurements; and inference for high-dimensional i.i.d. Gaussian vectors. Finally, in this paper we have not addressed the issue of model selection. A procedure for selection of M and r , based on an approximate leave-one-curve-out cross-validation score, has been proposed and implemented in Peng and Paul (2007). This approximation is based on a second order Taylor expansion of the negative log-likelihood at the estimator and it involves the intrinsic Fisher information operator and its inverse. Therefore, based on the analysis presented here, it is conjectured that the approximate CV score thus defined is asymptotically consistent for the class of models considered in this paper.

Appendix A : Properties of cubic B-spline basis

In many proofs of this paper, we need to use some properties of the cubic B -spline basis. We state some of them. More details can be found in de Boor (1978) and DeVore and Lorentz (1993). Let $\tilde{\phi} = (\tilde{\phi}_1, \dots, \tilde{\phi}_M)^T$ be the (standard) cubic B-spline basis functions on $[0, 1]$ with equally spaced knots. Then, the orthonormalized spline functions ϕ_1, \dots, ϕ_M are defined through $\phi(t) = G_{\phi, M}^{-1/2} \tilde{\phi}(t)$, where $G_{\phi, M} := ((\int \tilde{\phi}_k(t) \tilde{\phi}_l(t) dt))_{k, l=1}^M$, is the *Gram matrix* of $\tilde{\phi}$. It is known (cf. de Boor (1974), Burman (1985)) that $G_{\phi, M}$ is an $M \times M$ banded matrix, and satisfies,

$$\frac{c_{\phi, 0}}{M} I_M \leq G_{\phi, M} \leq \frac{c_{\phi, 1}}{M} I_M \quad \text{for some constants } 0 < c_{\phi, 0} < c_{\phi, 1} < \infty. \quad (29)$$

From this, and other properties of cubic B-splines (deVore and Lorentz, 1993, Chapter 13), we also have the following:

S1 $\sup_{t \in [0, 1]} \sum_{k=1}^M \phi_k^2(t) \leq c_{\phi, 2} M$ for some constant $c_{\phi, 2} > 0$.

S2 For any function $f \in C^{(4)}([0, 1])$, we have $\|f - P_{\phi, M}(f)\|_{\infty} = \|f^{(4)}\|_{\infty} O(M^{-4})$, where $P_{\phi, M}(f) = \sum_{k=1}^M \langle f, \phi_k \rangle \phi_k$ denotes the projection of f onto $\text{span}\{\phi_1, \dots, \phi_M\} = \text{span}\{\tilde{\phi}_1, \dots, \tilde{\phi}_M\}$.

Note that, property **S2** and assumption **A2** imply the existence of orthonormal functions $\{\psi_{*k}\}_{k=1}^r$ of the form

$$(\psi_{*1}(t), \dots, \psi_{*r}(t)) = (\psi_*(t))^T = B_*^T \phi(t), \quad B_*^T B_* = I_r,$$

which satisfy

$$\max_{1 \leq k \leq r} \|\bar{\psi}_k - \psi_{*k}\|_{\infty} \leq c_{\phi, 3} M^{-4} \max_{1 \leq k \leq r} \|\bar{\psi}_k^{(4)}\|_{\infty}. \quad (30)$$

Using these properties we obtain the following approximation to the important quantity $\|\Phi_i\|$, where $\Phi_i = [\phi(T_{i1}) : \dots : \phi(T_{im_i})]$ and $\|\cdot\|$ denotes the operator norm. This result will be extremely useful in the subsequent analysis.

Proposition 8: *Given $\eta > 0$, there is an event A_{η} defined in terms of the design points \mathbf{T} , with probability at least $1 - O(n^{-\eta})$, such that on the set A_{η} ,*

$$\|\Phi_i\|^2 \leq \bar{m} c_{g, 1} + \sqrt{5} c_{\phi, 0}^{-1} d_{\eta} [(M^{3/2} \log n) \vee (M \sqrt{\bar{m} \log n})], \quad (31)$$

for some constant $d_{\eta} > 0$, and $c_{g, 1}$ is the constant in condition **B2**. Furthermore, for all \mathbf{T} , we have the non-random bound

$$\|\Phi_i\|^2 \leq c_{\phi, 2} \bar{m} M, \quad \text{for all } i = 1, \dots, n. \quad (32)$$

Proof : First, (32) follows from the bound **S1**, since

$$\begin{aligned} \|\Phi_i\|^2 &= \|\Phi_i^T \Phi_i\| = \|\Phi_i \Phi_i^T\| \leq \|\Phi_i\|_F^2 = \text{tr}(\Phi_i^T \Phi_i) \\ &= \sum_{j=1}^{m_i} \sum_{k=1}^M (\phi_k(T_{ij}))^2 \leq c_{\phi, 2} m_i M \leq c_{\phi, 2} \bar{m} M, \quad \text{for all } i = 1, \dots, n. \end{aligned}$$

In order to prove (31), first write

$$\frac{1}{m_i} \Phi_i \Phi_i^T - \int \phi(t) (\phi(t))^T g(t) dt = G_{\phi, M}^{-1/2} \left[\left(\frac{1}{m_i} \sum_{j=1}^{m_i} [\tilde{\phi}_k(T_{ij}) \tilde{\phi}_l(T_{ij}) - \mathbb{E}(\tilde{\phi}_k(T_{ij}) \tilde{\phi}_l(T_{ij}))] \right)_{k, l=1}^M \right] G_{\phi, M}^{-1/2}. \quad (33)$$

Next, observe that, $\mathbb{E}[\phi_k(T_{i1})\phi_l(T_{i1})]^2 = \int (\tilde{\phi}_k(t))^2(\tilde{\phi}_l(t))^2 g(t) dt = 0$ for $|k - l| > 3$; and is within $[c_{\phi,4}c_{g,0}M^{-1}, c_{\phi,5}c_{g,1}M^{-1}]$, for constants $0 < c_{\phi,4} < c_{\phi,5} < \infty$, if $|k - l| \leq 3$. Then, using the fact that $\max_{1 \leq k \leq M} \|\phi_k\|_\infty$ is bounded, it follows from Bernstein's inequality that the set A_η defined by

$$A_\eta = \left\{ \mathbf{T} : \max_{1 \leq i \leq n} \max_{1 \leq k, l \leq M} \left| \frac{1}{m_i} \sum_{j=1}^{m_i} [\tilde{\phi}_k(T_{ij})\tilde{\phi}_l(T_{ij}) - \mathbb{E}(\tilde{\phi}_k(T_{ij})\tilde{\phi}_l(T_{ij}))] \right| \leq d_{1,\eta} \left(\frac{\log n}{m} \right) \vee \sqrt{\frac{\log n}{mM}} \right\} \quad (34)$$

has probability at least $1 - O(n^{-\eta})$, for some constant $d_{1,\eta} > 0$. Now, we can bound the Frobenius norm of the matrix in (33) by using (29) and (34), and the fact that the matrix has $O(M)$ nonzero elements. Then using (3) we derive (31).

Appendix B : Proofs for the sparse case

Proof of Proposition 2 : The main challenge in the proof of Proposition 2 is to efficiently approximate the average Kullback-Leibler divergence. We can express $K(\Sigma_i, \Sigma_{*i})$ as

$$K(\Sigma_i, \Sigma_{*i}) = \frac{1}{2} \sum_{j=1}^m [\lambda_j(R_{*i}) - \log(1 + \lambda_j(R_{*i}))], \quad (35)$$

where $\lambda_j(R_{*i})$ is the j -th largest eigenvalue of $R_{*i} = \Sigma_i^{-1/2}(\Sigma_{*i} - \Sigma_i)\Sigma_i^{-1/2}$. Using the inequality $e^x \geq 1 + x$ for $x \in \mathbb{R}$ (so that each term in the summation in (35) is nonnegative), and the Taylor series expansion for $\log(1 + x)$ for $|x| < 1$, it can be shown that, given $\epsilon > 0$ sufficiently small (but fixed), there exist constants $0 < c_{1,\epsilon} < c_{2,\epsilon} < \infty$ such that for $\|R_{*i}\|_F \leq \epsilon$,

$$c_{1,\epsilon} \|R_{*i}\|_F^2 \leq K(\Sigma_i, \Sigma_{*i}) \leq c_{2,\epsilon} \|R_{*i}\|_F^2. \quad (36)$$

Next, observe that

$$\frac{\|\Sigma_i^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F}{1 + \|\Sigma_{*i}^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F} \leq \|R_{*i}\|_F \leq \frac{\|\Sigma_{*i}^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F}{1 - \|\Sigma_{*i}^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F},$$

whenever $\|\Sigma_{*i}^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F < 1$. the proof of Proposition 1 can thus be reduced to finding probabilistic bounds for $\frac{1}{n} \sum_{i=1}^n \|\Sigma_{*i}^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F^2$.

One difficulty in obtaining those bounds is in handling the inverse of the matrices Σ_{*i} . In order to address that, and some related issues, we use the properties of the cubic spline basis derived in Appendix A. In the following lemmas we confine ourselves to the restricted parameter space $\Theta(\alpha_n)$, i.e., $(B, \Lambda) \in \Theta(\alpha_n)$.

Lemma 3: *Under the assumptions of Theorem 1 (for m_i 's bounded), or Corollary 2 (for m_i 's increasing slowly with n),*

$$(1 + d_1 \bar{\lambda}_1 r \bar{m})^{-1} \|\Sigma_i - \Sigma_{*i}\|_F \leq \|\Sigma_{*i}^{-1/2}(\Sigma_i - \Sigma_{*i})\Sigma_{*i}^{-1/2}\|_F \leq \|\Sigma_i - \Sigma_{*i}\|_F, \text{ for all } i = 1, \dots, n, \quad (37)$$

for some constant $d_1 > 0$

Proof : From condition **A2** and (30), it follows that, $\exists D_1 > 0$ such that, for all M ,

$$\max_{1 \leq k \leq r} \|\psi_{*k}\|_\infty \leq D_1 < \infty. \quad (38)$$

This, together with the definition of (B_*, Λ_*) , leads to the following bound on the eigenvalues of the matrices Σ_{*i} :

$$1 \leq \lambda_{\min}(\Sigma_{*i}) \leq \lambda_{\max}(\Sigma_{*i}) \leq 1 + D_1 r m_i \lambda_{*1} \leq 1 + d_1 \bar{\lambda}_1 r \bar{m}, \quad \text{for all } i = 1, \dots, n, \quad (39)$$

for some $d_1 > 0$, from which (37) follows.

Lemma 4: *Under the assumptions of Theorem 1 (for m_i 's bounded), or Corollary 2 (for m_i 's increasing slowly with n), given any $\eta > 0$, on the event A_η defined through (34) in Proposition 8, which has probability at least $1 - O(n^{-\eta})$, for sufficiently large n ,*

$$\max_{1 \leq i \leq n} \|\Sigma_i - \Sigma_{*i}\|_F^2 \leq \left[d_{3,\eta} \left(1 + d_1 \left[\left(\frac{M^{3/2} \log n}{\bar{m}} \right) \vee \sqrt{\frac{M^2 \log n}{\bar{m}}} \right] \bar{m}^2 \alpha_n^2 \right) \wedge [d_2 M \bar{m}^2 \alpha_n^2] \right], \quad (40)$$

where the second bound holds for all \mathbf{T} . Here $d_1, d_2, d_{3,\eta} > 0$ are appropriate constants depending on r and $\bar{\lambda}_1$.

Proof : An upper bound for $\|\Sigma_i - \Sigma_{*i}\|_F$ is obtained by expressing $I_M = B_* B_*^T + (I_M - B_* B_*^T)$, and then applying the triangle inequality,

$$\begin{aligned} & \|\Sigma_i - \Sigma_{*i}\|_F = \|\Phi_i^T (B \Lambda B^T - B_* \Lambda_* B_*^T) \Phi_i\|_F \\ & \leq \|\Phi_i^T B_* (B_*^T B \Lambda B^T B_* - \Lambda_*) B_*^T \Phi_i\|_F + 2 \|\Phi_i^T B_* B_*^T B \Lambda B^T (I_M - B_* B_*^T) \Phi_i\|_F \\ & \quad + \|\Phi_i^T (I_M - B_* B_*^T) B \Lambda B^T (I_M - B_* B_*^T) \Phi_i\|_F \\ & \leq \|\Phi_i^T B_*\|^2 \|B_*^T B \Lambda B^T B_* - \Lambda_*\|_F + 2 \|\Phi_i^T B_*\| \|\Lambda\| \|B^T (I_M - B_* B_*^T) \Phi\|_F \\ & \quad + \|\Lambda\| \|\Phi_i^T (I_M - B_* B_*^T) B\|_F^2 \\ & \leq D_1 r \bar{m} \|B_*^T B \Lambda B^T B_* - \Lambda_*\|_F \\ & \quad + \sqrt{d_4 r \bar{m} \bar{\lambda}_1} \|\Phi_i^T (I_M - B_* B_*^T) B\|_F \left[1 + (D_1 r \bar{m})^{-1/2} \|\Phi_i^T (I_M - B_* B_*^T) B\|_F \right], \quad (41) \end{aligned}$$

for some $d_4 > 1$. For the second inequality we use $\|B_*^T B\| \leq 1$, and for the last inequality we use (38) and (39). Next, by using (32), (17) and (5), we obtain the (nonrandom) bound

$$\bar{m}^{-1} \max_{1 \leq i \leq n} \|\Phi_i^T (I_M - B_* B_*^T) B\|_F^2 \leq c_{\phi,2} M \|(I_M - B_* B_*^T) B\|_F^2 \leq c_{\phi,2} M \alpha_n^2 (1 + o(1)) = o(1). \quad (42)$$

Then the bound in (41) can be majorized by,

$$D_1 r \bar{m} \|B_*^T B \Lambda B^T B_* - \Lambda_*\|_F + \sqrt{d_4 r \bar{m} \bar{\lambda}_1} \|\Phi_i\| \|(I_M - B_* B_*^T) B\|_F (1 + o(1)).$$

Using (31) to bound $\|\Phi_i\|$, from (41), and the definition of $\Theta(\alpha_n)$ together with (16) and (17), we obtain (40).

Lemma 5: *Under the assumptions of Theorem 1 (for m_i 's bounded), or Corollary 2 (for m_i 's increasing slowly with n), for any given $\eta > 0$, there is a positive sequence $\varepsilon_{1,n} = o(1)$ (depending on η), and a constant $d_{1,\eta} > 0$, such that*

$$\mathbb{P} \left[\frac{1}{n} \sum_{i=1}^n \|\Sigma_i - \Sigma_{*i}\|_F^2 > d_{1,\eta} \bar{m}^2 \alpha_n^2 (1 - \varepsilon_{1,n}) \right] \geq 1 - O(n^{-(2+2\kappa)Mr-\eta}),$$

where κ is as in **B1**.

Proof : Let $\Delta = B\Lambda B^T - B_*\Lambda_*B_*^T$. Observe that, by definition of $\Theta(\alpha_n)$ (equation (15)), and equations (16) and (17), for large n ,

$$\|\Delta\|_F^2 \leq c_*\alpha_n^2, \quad (\text{for some constant } c_* > 0). \quad (43)$$

First, consider the lower bound

$$\|\Sigma_i - \Sigma_{*i}\|_F^2 = \text{tr}[\Phi_i^T \Delta \Phi_i \Phi_i^T \Delta \Phi_i] \geq \sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2. \quad (44)$$

We are going to derive an exponential tail bound for $\frac{1}{n} \sum_{i=1}^n \sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2$. Rewriting the term on the extreme right, and using the fact that $\{T_{ij}\}_{j=1}^{m_i}$ are i.i.d. with density $g : c_{g,0} \leq g \leq c_{g,1}$ (**B2**), we have, for all i ,

$$\begin{aligned} \mathbb{E} \sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2 &= m_i(m_i - 1) \text{tr}(\mathbb{E}[\phi(T_{i1})(\phi(T_{i1}))^T] \Delta \mathbb{E}[\phi(T_{i2})(\phi(T_{i2}))^T] \Delta) \\ &\in (c_{g,0}^2 \underline{m}(\underline{m} - 1) \|\Delta\|_F^2, c_{g,1}^2 \overline{m}(\overline{m} - 1) \|\Delta\|_F^2) \\ &\in (d'_1 \underline{m}^2 \alpha_n^2 (1 + o(1)), d''_1 \overline{m}^2 \alpha_n^2 (1 + o(1))), \end{aligned} \quad (45)$$

for some $d''_1 \geq d'_1 > 0$ (whose values depend on $c_{g,0}$, $c_{g,1}$ and the constants appearing in **A1**), where in the last step we use (43). The last inequality uses (72), (73), the definition of $\Theta(\alpha_n)$, and the properties (16) and (17). Notice that, the variance of $\sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2$ can be bounded, for sufficiently large n , as

$$\begin{aligned} &\max_{1 \leq i \leq n} \text{Var} \left(\sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2 \right) \\ &\leq \max_{1 \leq i \leq n} \mathbb{E} \left(\|\Sigma_i - \Sigma_{*i}\|_F^2 \sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2 \right) \\ &\leq d'_2 M(\overline{m}^2 \alpha_n^2)^2 =: V_{1,n}, \end{aligned}$$

where $d'_2 > 0$ is some constant. In the above, we obtain the first inequality by (44), and the second inequality by using (40) and (45). Next, $\sum_{j_1 \neq j_2}^{m_i} [(\phi(T_{ij_1}))^T \Delta \phi(T_{ij_2})]^2$, for $i = 1, \dots, n$, are independent, and bounded by $K_{1,n} := d_4 M \overline{m}^2 \alpha_n^2$, for a constant $d_4 > 0$ (using (40)). Hence, by applying Bernstein's inequality, and noticing that $K_{1,n} \sqrt{\frac{M \log n}{n}} = o(\sqrt{V_{1,n}})$, and $\sqrt{V_{1,n}} \sqrt{\frac{M \log n}{n}} = o(\overline{m}^2 \alpha_n^2)$ (by (5), or (6)), the result follows.

Lemma 6: *Under the assumptions of Theorem 1 (for m_i 's bounded), or Corollary 2 (for m_i 's increasing slowly with n), for any given $\eta > 0$, there is a positive sequence $\varepsilon_{2,n} = o(1)$ and a constant $d_{2,\eta} > 0$, such that*

$$\mathbb{P} \left[\frac{1}{n} \sum_{i=1}^n \|\Sigma_i - \Sigma_{*i}\|_F^2 < d_{2,\eta} \overline{m}^2 \alpha_n^2 (1 + \varepsilon_{2,n}) \right] \geq 1 - O(n^{-(2+2\kappa)Mr-\eta}), \quad (46)$$

where κ is as in **B1**.

Proof : From the proof of Lemma 4, especially the inequalities (41) and (42), it is clear that we only need to provide a sharp upper bound for $\frac{1}{n} \sum_{i=1}^n \|\Phi_i^T (I_M - B_* B_*^T) B\|_F^2$. Let $\overline{\Delta} := (I_M - B_* B_*^T) B$. Then from (17), for n large enough,

$$\|\overline{\Delta}\|_F^2 \leq c_* \alpha_n^2 \quad (47)$$

for some $c_* > 0$. Then, using (3), for all i ,

$$\begin{aligned} \mathbb{E} \|\Phi_i^T (I_M - B_* B_*^T) B\|_F^2 &= \sum_{j=1}^{m_i} \text{tr} \left(\mathbb{E}(\phi(T_{ij})(\phi(T_{ij}))^T) \overline{\Delta} \overline{\Delta}^T \right) \\ &\leq c_{g,1} m_i \text{tr} [\overline{\Delta} \overline{\Delta}^T] \leq c_{g,1} \overline{m} \|\overline{\Delta}\|_F^2. \end{aligned} \quad (48)$$

Combining (48) with (47), (41) and (42), we get, for sufficiently large n , and some constant $C > 0$,

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} \|\Sigma_i - \Sigma_{*i}\|_F^2 \leq C \overline{m}^2 \alpha_n^2$$

Next, using (42), we have

$$\begin{aligned} \max_{1 \leq i \leq n} \text{Var}(\|\Phi_i^T \overline{\Delta}\|_F^2) &\leq \max_{1 \leq i \leq n} \mathbb{E} \|\Phi_i^T \overline{\Delta}\|_F^4 \\ &\leq c_{\phi,2} \overline{m} M \|\overline{\Delta}\|_F^2 \mathbb{E} \|\Phi_i^T \overline{\Delta}\|_F^2 \\ &\leq c_{\phi,2} c_{g,1} \overline{m}^2 M \|\overline{\Delta}\|_F^4 \\ &\leq C' M \overline{m}^2 \alpha_n^4 (1 + \varepsilon_n) =: V_{2,n}, \end{aligned} \quad (49)$$

for some positive sequence $\varepsilon_n = o(1)$ and some constant $C' > 0$, where in the last step we used (47). Again, using Bernstein's inequality for $\frac{1}{n} \sum_{i=1}^n \|\Phi_i^T \overline{\Delta}\|_F^2$, which is a sum of independent variables bounded by $K_{2,n} = c_{g,1} M \overline{m} \alpha_n^2 (1 + o(1))$, the result follows (checking that, by (5) or (6), we have, $K_{2,n} \sqrt{\frac{M \log n}{n}} = o(\sqrt{V_{2,n}})$ and $\sqrt{V_{2,n}} \sqrt{\frac{M \log n}{n}} = o(\overline{m}^2 \alpha_n^2)$).

Proof of Proposition 3 : Write, $R_i = \overline{\Sigma}_i^{1/2} (\Sigma_i^{-1} - \Sigma_{*i}^{-1}) \overline{\Sigma}_i^{1/2}$. We can bound $\|R_i\|_F$ as

$$\begin{aligned} \|R_i\|_F &\leq \|\overline{\Sigma}_i^{1/2} \Sigma_i^{-1/2}\| \|\overline{\Sigma}_i^{1/2} \Sigma_{*i}^{-1/2}\| \|\Sigma_i^{-1/2} \Sigma_{*i}^{1/2}\| \|\Sigma_{*i}^{-1/2} (\Sigma_i - \Sigma_{*i}) \Sigma_{*i}^{-1/2}\|_F \\ &\leq \|\overline{\Sigma}_i^{1/2} \Sigma_{*i}^{-1/2}\|^2 \|\Sigma_i^{-1/2} \Sigma_{*i}^{1/2}\|^2 \|\Sigma_i - \Sigma_{*i}\|_F \\ &\leq (1 + \|\overline{\Sigma}_i - \Sigma_{*i}\|) (1 - \|\Sigma_i - \Sigma_{*i}\|)^{-1} \|\Sigma_i - \Sigma_{*i}\|_F, \end{aligned}$$

where the third inequality is due to (39). Note that, by condition **C**, it follows that $\max_{1 \leq i \leq n} \|\overline{\Sigma}_i - \Sigma_{*i}\|_F \leq C \overline{m} \overline{\beta}_n = o(1)$ for some constant $C > 0$. Therefore, applying (40), we observe that for $\mathbf{T} \in A_\eta$ with A_η as in (34), for large enough n , $\|R_i\|_F \leq 2 \|\Sigma_i - \Sigma_{*i}\|_F$. Due to the Gaussianity of the observations, for any symmetric $m_i \times m_i$ matrix A , the random variable $\text{tr}(A(S_i - \overline{\Sigma}_i))$ has the same distribution as $\text{tr}(D_i(X_i X_i^T - I_{m_i}))$, where D_i is a diagonal matrix of the eigenvalues of $\overline{\Sigma}_i^{-1/2} A \overline{\Sigma}_i^{-1/2}$, and $X_i \sim N(0, I_{m_i})$ are independent. Therefore, using an exponential inequality for a weighted sum of independent χ_1^2 random variables, we have, for $\mathbf{T} \in A_\eta$ and each $(B, \Lambda) \in \Theta(\alpha_n)$,

$$\begin{aligned} \mathbb{P} \left(\left| \frac{1}{n} \sum_{i=1}^n \text{tr}((\Sigma_i^{-1} - \Sigma_{*i}^{-1})(S_i - \overline{\Sigma}_i)) \right| \leq d_{3,\eta} \sqrt{\frac{M \log n}{n}} \left(\frac{1}{n} \sum_{i=1}^n \|\Sigma_i - \Sigma_{*i}\|_F^2 \right)^{1/2} \mid \mathbf{T} \right) \\ \geq 1 - O(n^{-(2+2\kappa)Mr-\eta}), \end{aligned}$$

for a constant $d_{3,\eta} > 0$. Therefore, using (46) we conclude the proof.

Proof of Proposition 4 : Using Cauchy-Schwarz inequality twice, we can bound the last term in (19), which corresponds to *model bias*, as

$$\begin{aligned} \left| \frac{1}{2n} \sum_{i=1}^n \text{tr}((\Sigma_i^{-1} - \Sigma_{*i}^{-1})(\bar{\Sigma}_i - \Sigma_{*i})) \right| &\leq \frac{1}{2} \left[\frac{1}{n} \sum_{i=1}^n \|\bar{\Sigma}_i - \Sigma_{*i}\|_F^2 \right]^{1/2} \left[\frac{1}{n} \sum_{i=1}^n \|\Sigma_i^{-1} - \Sigma_{*i}^{-1}\|_F^2 \right]^{1/2} \\ &\leq \frac{1}{2} \max_{1 \leq j \leq n} \|\bar{\Sigma}_j - \Sigma_{*j}\|_F \left[\frac{1}{n} \sum_{i=1}^n \|\Sigma_i^{-1} - \Sigma_{*i}^{-1}\|_F^2 \right]^{1/2} \\ &\leq \frac{1}{2} \bar{m} \bar{\beta}_n \max_{1 \leq j \leq n} \|\Sigma_j^{-1}\| \|\Sigma_{*j}^{-1}\| \left[\frac{1}{n} \sum_{i=1}^n \|\Sigma_i - \Sigma_{*i}\|_F^2 \right]^{1/2}, \end{aligned}$$

where in the last step we used (71). Thus, the proof is finished by using condition **C**, (39) and (46).

Now, to complete the proof of Theorem 1, we give the details of the covering argument.

Proof of Lemma 2 : Using an expansion analogous to (19) and the upper bound in (36), and applying Cauchy-Schwarz inequality, we have, for some constants $C_1, C_2 > 0$, on $A_{4,\eta}$ and for $\mathbf{T} \in A_\eta$, for n large enough,

$$\begin{aligned} &|L_n(B_1, \Lambda_1) - L_n(B_2, \Lambda_2)| \\ &\leq \frac{1}{n} \sum_{i=1}^n \left[\|\Sigma_{1,i}^{-1} - \Sigma_{2,i}^{-1}\|_F \|S_i - \bar{\Sigma}_i\|_F + C_1 \|\Sigma_{1,i} - \Sigma_{2,i}\|_F^2 + \|\Sigma_{1,i}^{-1} - \Sigma_{2,i}^{-1}\|_F \|\bar{\Sigma}_i - \Sigma_{2,i}\|_F \right] \\ &\leq \max_{1 \leq i \leq n} \|\Sigma_{1,i}^{-1} - \Sigma_{2,i}^{-1}\|_F \|\bar{\Sigma}_i\| \max_{1 \leq i \leq n} \|\bar{\Sigma}_i^{-1/2} S_i \bar{\Sigma}_i^{-1/2} - I_{m_i}\|_F + C_2 \max_{1 \leq i \leq n} \|\Sigma_{1,i} - \Sigma_{2,i}\|_F^2 \\ &\quad + \max_{1 \leq i \leq n} \|\Sigma_{1,i}^{-1} - \Sigma_{2,i}^{-1}\|_F \left(\max_{1 \leq i \leq n} \|\bar{\Sigma}_i - \Sigma_{*i}\|_F + \max_{1 \leq i \leq n} \|\Sigma_{2,i} - \Sigma_{*i}\|_F \right) \\ &\leq d_{4,\eta} [\bar{m}^2 \delta_n \bar{m} \log n + \bar{m}^2 \delta_n^2 + \bar{m} \delta_n (\bar{m} \bar{\beta}_n + \sqrt{M \bar{m} \alpha_n})] = o(\alpha_n^2). \end{aligned}$$

In the last step, we have used Lemma 4 (for the last term), the identity (71) in Appendix F, and the fact that $\|\Sigma_{k,i}^{-1}\| \leq 1$ ($k = 1, 2$).

Proof of Corollary 1: The best rate follows by direct calculation.

The near-optimality of the estimator requires proving that, for an appropriately chosen subclass \mathcal{C} of covariance kernels of rank r , we have the following analog of Theorem 2 of Hall *et al.* (2006): for any estimator $\{\hat{\psi}_k\}_{k=1}^r$ of the eigenfunctions $\{\psi_k\}_{k=1}^r$, for n sufficiently large,

$$\min_{1 \leq k \leq r} \sup_{\bar{\Sigma}_0 \in \mathcal{C}} \mathbb{E} \|\hat{\psi}_k - \bar{\psi}_k\|_2^2 \geq C n^{-8/9}, \quad (50)$$

for some $C > 0$. Here the parameter space \mathcal{C} consists of covariance kernels of rank r with eigenfunctions satisfying **A1-A2**. Moreover, the random design satisfies **B1-B2**, with \bar{m} bounded above.

The derivation of the lower bound on the risk involves construction of a finite, “least favorable” parameter set in \mathcal{C} by combining the constructions in Paul and Johnstone (2007) (for obtaining lower bounds on risk in high-dimensional PCA) and Hall *et al.* (2006) (for functional data case). This construction is as follows. Let $\phi_1^0, \dots, \phi_r^0$ be a set of orthonormal functions on $[0, 1]$ which

are four times continuously differentiable, with fourth derivative bounded. Let $M_* \asymp n^{1/9}$ be an integer appropriately chosen. Let $\gamma_1, \dots, \gamma_{M_*}$ be a set of basis functions that are (i) orthonormal on $[0, 1]$, and orthogonal to the set $\{\phi_1^0, \dots, \phi_r^0\}$; (ii) are four times continuously differentiable and γ_j is supported on an interval of length $O(M_*^{-1})$ around the point $\frac{j}{M_*}$. One particular choice for these functions is to let $\{\phi_k^0\}$ be the translated periodized scaling functions of a wavelet basis at a certain *scale* with adequate degree of smoothness, and to let $\{\gamma_j\}_{j=1}^{M_*}$ be the set of compactly supported, orthonormal, periodized wavelet functions corresponding to the scaling functions. Indeed, then we can choose M_* to be an integer power of 2. Note that, such a basis $(\{\phi_k^0 : 1 \leq k \leq r\} \cup \{\gamma_l : 1 \leq l \leq M_*\})$ has the stability and smoothness property commensurate with the orthonormalized B-spline basis we are using for deriving the REML estimators. Next, let $\bar{\lambda}_1 > \dots > \bar{\lambda}_r > 0$ be fixed numbers satisfying **A1**. Finally, let us define a covariance kernel $\bar{\Sigma}_0^{(0)}$ as

$$\bar{\Sigma}_0^{(0)}(s, t) = \sum_{k=1}^r \bar{\lambda}_k \phi_k^0(s) \phi_k^0(t), \quad s, t \in [0, 1]. \quad (51)$$

Also, for each fixed j in some index set \mathcal{F}_0 (to be specified below), define

$$[\psi_1^{(j)}(s) : \dots : \psi_r^{(j)}(s)] = \tilde{\Psi}(s) \bar{B}^{(j)}, \quad s \in [0, 1]$$

where $\tilde{\Psi}(s) = (\phi_1^0(s), \dots, \phi_r^0(s), \gamma_1(s), \dots, \gamma_{M_*}(s))$ and $\bar{B}^{(j)}$ is an $(M_* + r) \times r$ matrix with orthonormal columns (to be specified below). Then define

$$\bar{\Sigma}_0^{(j)}(s, t) = \sum_{k=1}^r \bar{\lambda}_k \psi_k^{(j)}(s) \psi_k^{(j)}(t), \quad s, t \in [0, 1], \quad (52)$$

for $j \in \mathcal{F}_0$. We require that $\log |\mathcal{F}_0| \asymp M_* \asymp n^{1/9}$, and $\|\bar{B}^{(j)} - \bar{B}^{(j')}\|_F^2 \asymp n^{-8/9}$, for $j \neq j'$ and $j, j' \in \mathcal{F}_0 \cup \{0\}$. Here $\bar{B}^{(0)}$ is the $(M_* + r) \times r$ matrix of basis coefficients of $\bar{\Sigma}_0^{(0)}$ with columns \mathbf{e}_k , the k -th canonical basis vector in \mathbb{R}^{M_*+r} .

The proof of the minimax lower bound is based on an application of *Fano's Lemma* (Yang and Barron, 1999), which requires computation of the Kullback-Leibler divergence between two specific values of the parameters. In order to apply *Fano's lemma*, we need to choose \mathcal{F}_0 and $\bar{B}^{(j)}, j \in \mathcal{F}_0$, such that

$$\frac{\text{ave}_{j \in \mathcal{F}_0} [\sum_{i=1}^n \mathbb{E} K(\bar{\Sigma}_i^{(j)}, \bar{\Sigma}_i^{(0)})] + \log 2}{\log |\mathcal{F}_0|} \approx c \in (0, 1), \quad (53)$$

where $\bar{\Sigma}_i^{(j)}$ denotes the covariance of the observation i given $\{T_{il}\}_{l=1}^{m_i}$ under the model parameterized by $\bar{\Sigma}_0^{(j)}$, and \mathbb{E} denotes expectation with respect to the design points \mathbf{T} . Under the assumptions on the design points, using the properties of the basis functions $\{\phi_k^0\}_{k=1}^r$ and $\{\gamma_k\}_{k=1}^{M_*}$, and the computations carried out in the proof of Proposition 2 (in Appendix B), in particular a nonrandom bound analogous to the second bound appearing in Lemma 4, it is easy to see that for n large enough (so that $\|\bar{B}^{(j)} - \bar{B}^{(0)}\|_F$ is sufficiently small), we have

$$\frac{1}{n} \sum_{i=1}^n K(\bar{\Sigma}_i^{(j)}, \bar{\Sigma}_i^{(0)}) \asymp \frac{1}{n} \sum_{i=1}^n \|\bar{\Sigma}_i^{(j)} - \bar{\Sigma}_i^{(0)}\|_F^2.$$

From this, and the property of the basis used to represent the eigenfunctions, it follows that

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} K(\bar{\Sigma}_i^{(j)}, \bar{\Sigma}_i^{(0)}) \asymp \| \bar{B}^{(j)} - \bar{B}^{(0)} \|_F^2. \quad (54)$$

The task remains to construct \mathcal{F}_0 and $\bar{B}^{(j)}$ appropriately so that $\mathcal{C}_0 := \{\bar{\Sigma}_0^{(j)} : j \in \mathcal{F}_0\}$ is in \mathcal{C} , for n sufficiently large.

Following the proof of Theorem 2 in Paul and Johnstone (2007), we first define $M_0 = \lfloor \frac{2M_*}{9r} \rfloor$. Then define the k -th column of $\bar{B}^{(j)}$ as

$$\bar{B}_k^{(j)} = \sqrt{1 - \delta_k^2} \mathbf{e}_k + \delta_k \sum_{l=1}^{M_*} z_{kl}^{(j)} \mathbf{e}_{r+l}, \quad k = 1, \dots, r, \quad (55)$$

where $z_{kl}^{(j)}$ are appropriately chosen using a “sphere packing” argument (to ensure that $\log |\mathcal{F}_0| \asymp M_*$), and take values in $\{-M_0^{-1/2}, 0, M_0^{-1/2}\}$. Moreover, let S_k be the set of coordinates l such that $z_{kl}^{(j)} \neq 0$ for some $j \in \mathcal{F}_0$. By construction, S_k are disjoint for different $k = 1, \dots, r$, and $|S_k| \sim M_*/r$. Hence,

$$\psi_k^{(j)} = \sqrt{1 - \delta_k^2} \phi_k^0 + \delta_k \sum_{l \in S_k} z_{kl}^{(j)} \gamma_l, \quad k = 1, \dots, r. \quad (56)$$

Furthermore, by the construction of $\{z_{lk}^{(j)}\}$, $\sum_{l \in S_k} |z_{lk}^{(j)}|^2 = 1$, and for any $j \neq j'$ the vectors $\mathbf{z}_k^{(j)} = (z_{kl}^{(j)})_{l \in S_k}$ and $\mathbf{z}_k^{(j')} = (z_{kl}^{(j')})_{l \in S_k}$ satisfy $\|\mathbf{z}_k^{(j)} - \mathbf{z}_k^{(j')}\|_2 \geq 1$. Therefore, from (55) it follows that the RHS of (54) is of the order δ_k^2 , and hence in order that (53) is satisfied, we need to choose $\delta_k \sim n^{-4/9} \asymp M_*^{-4}$. It follows immediately from (56) that (i) the eigenfunctions $\psi_1^{(j)}, \dots, \psi_r^{(j)}$ are orthonormal, and four times continuously differentiable. Also, since γ_l is centered around l/M_* with a support of the order $O(M_*^{-1})$, it follows that, for only finitely many $l \neq l'$, the support of $\gamma_{l'}$ overlaps with the support of γ_l . Moreover, if $\gamma_l^{(s)}$ denotes the s -th derivative of γ_l , then $\|\gamma_l^{(s)}\|_\infty = O(M_*^{1/2+s})$, for $s = 0, 1, \dots, 4$. Thus, the choice $\delta_k \asymp M_*^{-4}$ ensures that, (ii) for each $k = 1, \dots, r$, the fourth derivative of $\psi_k^{(j)}$ is bounded. Hence, by appropriate choice of the constants, we have that $\mathcal{C}_0 \subset \mathcal{C}$. Finally, arguing as in Paul and Johnstone (2007), with an application of the *Fano's Lemma* we conclude (50).

Appendix C : Proof of Proposition 5

Using standard arguments, it can be shown that, given $\epsilon > 0$ sufficiently small (but fixed), we have constants $0 < c_{1,\epsilon} < c_{2,\epsilon} < \infty$ such that for $\|\Gamma_*^{-1/2}(\Gamma_* - \Gamma)\Gamma_*^{-1/2}\|_F \leq \epsilon$,

$$c_{1,\epsilon} \|\Gamma_*^{-1/2}(\Gamma_* - \Gamma)\Gamma_*^{-1/2}\|_F^2 \leq K(\Gamma, \Gamma_*) \leq c_{2,\epsilon} \|\Gamma_*^{-1/2}(\Gamma_* - \Gamma)\Gamma_*^{-1/2}\|_F^2. \quad (57)$$

Thus, it suffices to provide tight bounds for $\|\Gamma_*^{-1/2}(\Gamma - \Gamma_*)\Gamma_*^{-1/2}\|_F$. We introduce some notations first. Define, $G = \frac{\sigma^2}{s} \Lambda^{-1} + I_r$, $G_* = \frac{\sigma^2}{s} \Lambda_*^{-1} + I_r$ and $\Delta = B \Lambda B^T - B_* \Lambda_* B_*^T$. Then,

$$\Gamma^{-1} = \frac{1}{\sigma^2} (I_M - B(\frac{\sigma^2}{s} \Lambda^{-1} + I_r)^{-1} B^T) = \frac{1}{\sigma^2} (I_M - B G^{-1} B^T), \quad \text{and} \quad \Gamma_*^{-1} = \frac{1}{\sigma^2} (I_M - B_* G_*^{-1} B_*^T).$$

Moreover, due to **A1'**, there exist constants, $c_3, c_4 > 0$, such that,

$$c_3 \left(\frac{\sigma^2}{s} \right) \leq \sigma_{\min}(I_r - G_*^{-1}) \leq \sigma_{\max}(I_r - G_*^{-1}) \leq c_4 \left(\frac{\sigma^2}{s} \right).$$

We express $I_M - B_* G_*^{-1} B_*^T$ as $(I_M - B_* B_*^T) + B_*(I_r - G_*^{-1})B_*^T$. Then we can express $(\sigma^2/s)^2 \|\Gamma_*^{-1/2}(\Gamma - \Gamma_*)\Gamma_*^{-1/2}\|_F^2 = \sigma^4 \|\Gamma_*^{-1/2}\Delta\Gamma_*^{-1/2}\|_F^2$ as

$$\begin{aligned} & \text{tr}[(I_M - B_* G_*^{-1} B_*^T)\Delta(I_M - B_* G_*^{-1} B_*^T)\Delta] \\ &= \text{tr}[(I_M - B_* B_*^T)B\Lambda B^T(I_M - B_* B_*^T)B\Lambda B^T] + 2\text{tr}[B_*(I_r - G_*^{-1})B_*^T B\Lambda B^T(I_M - B_* B_*^T)B\Lambda B^T] \\ & \quad + \text{tr}[B_*(I_r - G_*^{-1})B_*^T \Delta B_*(I_r - G_*^{-1})B_*^T \Delta] \\ &= \|(I_M - B_* B_*^T)B\Lambda B^T(I_M - B_* B_*^T)\|_F^2 \\ & \quad + 2\text{tr}[(I_r - G_*^{-1})^{1/2}B_*^T B\Lambda(B^T(I_M - B_* B_*^T)B)\Lambda B^T B_*(I_r - G_*^{-1})^{1/2}] \\ & \quad + \|(I_r - G_*^{-1})^{1/2}(B_*^T B\Lambda B^T B_* - \Lambda_*)(I_r - G_*^{-1})^{1/2}\|_F^2 \\ &\geq 2c_3\lambda_r^2(\sigma_{\min}(B_*^T B))^2 \frac{\sigma^2}{s} \|(I_M - B_* B_*^T)B\|_F^2 + c_3^2 \left(\frac{\sigma^2}{s} \right)^2 \|B_*^T B\Lambda B^T B_* - \Lambda_*\|_F^2 \\ &\geq c_4 \left(\frac{\sigma^2}{s} \right) \|(I_M - B_* B_*^T)B\|_F^2 + c_3^2 \left(\frac{\sigma^2}{s} \right)^2 \|B_*^T B\Lambda B^T B_* - \Lambda_*\|_F^2 \end{aligned} \quad (58)$$

for constants $c_3, c_4 > 0$. Now, since $(B, \Lambda) \in \tilde{\Theta}(\alpha_n)$, where $B = \mathbf{exp}(1, B_* A_U + C_U)$ it follows that $\|A_U\|_F \leq \alpha_n \sqrt{\frac{s}{\sigma^2}}$ and $\|C_U\|_F \leq \alpha_n$. Moreover, from (16), and using the fact that $A_U = -A_U^T$, we have,

$$B_*^T B\Lambda B^T B_* - \Lambda_* = D + (A_U \Lambda - \Lambda A_U) + O(\|A_U\|_F^2 + \|D\|_F^2 + \|U\|_F(\|A_U\|_F + \|C_U\|_F)).$$

Since $A_U \Lambda - \Lambda A_U$ is symmetric, has zeros on the diagonal, and its Frobenius norm is bounded below by $\min_{1 \leq j < k \leq r}(\lambda_j - \lambda_k) \|A_U\|_F$, and D is diagonal, it follows that for some constant $c_6 > 0$,

$$\|B_*^T B\Lambda B^T B_* - \Lambda_*\|_F^2 \geq c_6(\|D\|_F^2 + \|A_U\|_F^2) - O\left(\left(\frac{s}{\sigma^2}\right)^{3/2} \alpha_n^3\right).$$

From this, and using (17) to approximate the first term in (58), it follows that for some constant $c_7 > 0$,

$$\begin{aligned} & \|(I_M - B_* G_*^{-1} B_*^T)^{1/2} \Delta (I_M - B_* G_*^{-1} B_*^T)^{1/2}\|_F^2 \\ &\geq c_7 \left(\frac{\sigma^2}{s} \right) \left[\|C_U\|_F^2 + \frac{\sigma^2}{s} \|A_U\|_F^2 + \frac{\sigma^2}{s} \|D\|_F^2 - O\left(\sqrt{\frac{s}{\sigma^2}} \alpha_n^3\right) \right] \\ &= c_7 \left(\frac{\sigma^2}{s} \right) \alpha_n^2 (1 - o(1)). \end{aligned} \quad (59)$$

The last equality is because $\alpha_n \sqrt{\frac{s}{\sigma^2}} = o(1)$. Also, it is easy to show now that, for some $c_8 > 0$

$$\|(I_M - B_* G_*^{-1} B_*^T)^{1/2} \Delta (I_M - B_* G_*^{-1} B_*^T)^{1/2}\|_F^2 \leq c_8 \left(\frac{\sigma^2}{s} \right) \alpha_n^2 (1 + o(1)). \quad (60)$$

Hence, from (59) and (60), it follows that, there are constants $c_9, c_{10} > 0$ such that, for sufficiently large n ,

$$c_9 \alpha_n \sqrt{\frac{s}{\sigma^2}} \leq \|\Gamma_*^{-1/2}(\Gamma - \Gamma_*)\Gamma_*^{-1/2}\|_F \leq c_{10} \alpha_n \sqrt{\frac{s}{\sigma^2}},$$

which, together with (57), proves (23).

Appendix D : Score representation in the matrix case

First, define the *canonical metric* on the tangent space $\mathcal{T}_B \oplus \mathbb{R}^r$ of the parameter space $\tilde{\Omega} = \mathcal{S}_{M,r} \otimes \mathbb{R}^r$ (for $\theta = (B, \zeta)$) by

$$\langle X, Y \rangle_g = \langle X_B, Y_B \rangle_c + \langle X_\zeta, Y_\zeta \rangle, \quad \text{for } X_B, Y_B \in \mathcal{T}_B, \quad X_\zeta, Y_\zeta \in \mathbb{R}^r,$$

where $\langle X_B, Y_B \rangle_c = \text{tr}(X_B^T(I_M - \frac{1}{2}BB^T)Y_B)$ is the canonical metric on $\mathcal{S}_{M,r}$ and $\langle X_\zeta, Y_\zeta \rangle = \text{tr}(X_\zeta^T Y_\zeta)$ is the usual Euclidean metric. Next, for an arbitrary θ , write $\tilde{L}_n(\theta) = F_n^1(\theta) + F_n^2(\theta)$, where $F_n^1(\theta) = \text{tr}(\Gamma^{-1}\tilde{S})$ and $F_n^2(\theta) = \log|\Gamma| = \log|I_r + e^\zeta| = \log|I_r + \Lambda|$. Similarly, we write $L(\theta; \theta_*) = F^1(\theta; \theta_*) + F^2(\theta; \theta_*)$, where $F^1(\theta; \theta_*) = \text{tr}(\Gamma^{-1}\Gamma_*)$ and $F^2(\theta; \theta_*) = F_n^2(\theta)$. Below, we shall only give expressions for gradient and Hessian of $F_n^1(\cdot)$ and $F_n^2(\cdot)$, since the gradient and Hessian of $F^1(\cdot; \theta_*)$ and $F^2(\cdot; \theta_*)$ follow from these (by replacing \tilde{S} with Γ_*).

Gradient and Hessian

From Appendices B and D of Peng and Paul (2007), we obtain expressions for the gradient and Hessian of $F_n^1(\cdot)$ and $F_n^2(\cdot)$. We mainly follow the notations used there. Let, $P := P(\theta) = I_M + B\Lambda B^T$. Then $P^{-1} = I_M - BQ^{-1}B^T$, where

$$Q := Q(\theta) = \Lambda^{-1} + B^T B = \Lambda^{-1} + I_r \implies Q^{-1} = \Lambda(I_r + \Lambda)^{-1}.$$

The fact that Q is independent of B is of importance in the calculations throughout. Use $F_{n,B}^1(\cdot)$ to denote the Euclidean gradient of $F_n^1(\cdot)$ w.r.t. B . It is easy to see that $F_{n,B}^1(\theta) = -2\tilde{S}BQ^{-1}$. Then, under the canonical metric the intrinsic gradient is given by

$$\nabla_B F_n^1(\theta) = F_{n,B}^1(\theta) - B(F_{n,B}^1(\theta))^T B = 2[BQ^{-1}B^T\tilde{S}B - \tilde{S}BQ^{-1}].$$

Since $F_n^2(\theta)$ does not involve B , the Euclidean gradient $F_{n,B}^2(\theta) = 0$, and hence $\nabla_B F_n^2(\theta) = 0$. Therefore,

$$\nabla_B \tilde{L}_n(\theta) = \nabla_B F_n^1(\theta) = 2[BQ^{-1}B^T\tilde{S}B - \tilde{S}BQ^{-1}]. \quad (61)$$

Next, for $X_B \in \mathcal{T}_B$, let $G_{n,BB}^1(\cdot)(X_B)$ be the Euclidean Hessian operator of $F_n^1(\cdot)$ evaluated at X_B . It is computed as

$$G_{n,BB}^1(\theta)(X_B) = -2\tilde{S}X_BQ^{-1}.$$

The Hessian operator of $\tilde{L}_n(\cdot)$ w.r.t. B , equals the Hessian operator of $F_n^1(\cdot)$ w.r.t. B . For $X_B, Y_B \in \mathcal{T}_B$, it is given by

$$\begin{aligned} H_{n,B}(\theta)(X_B, Y_B) &= \text{tr}(Y_B^T G_{n,BB}^1(\theta)(X_B)) + \frac{1}{2} \text{tr}([(F_{n,B}^1(\theta))^T X_B B^T + B^T X_B (F_{n,B}^1(\theta))^T] Y_B) \\ &\quad - \frac{1}{2} \text{tr}[(B^T F_{n,B}^1(\theta) + (F_{n,B}^1(\theta))^T B) X_B^T (I_M - BB^T) Y_B]. \end{aligned} \quad (62)$$

For computing gradient and Hessian with respect to ζ , we only need to compute first and second derivatives of the function $\tilde{L}_n(\cdot)$ (equivalently, of $F_n^1(\cdot)$ and $F_n^2(\cdot)$). Using calculations carried out in Appendix D of Peng and Paul (2007), and the identity $P^{-1}B_k = (1 + \lambda_k)^{-1}B_k$ where B_k is the k -th column of B , $1 \leq k \leq r$, we have,

$$\frac{\partial F_n^1}{\partial \zeta_k}(\theta) = -e^{\zeta_k} B_k^T P^{-1} \tilde{S} P^{-1} B_k = -\frac{\lambda_k}{(1 + \lambda_k)^2} B_k^T \tilde{S} B_k, \quad \text{and} \quad \frac{\partial F_n^2}{\partial \zeta_k}(\theta) = e^{\zeta_k} B_k^T P^{-1} B_k = \frac{\lambda_k}{1 + \lambda_k}. \quad (63)$$

Thus

$$\nabla_{\zeta} \tilde{L}_n(\theta) = \text{diag} \left(\frac{\lambda_k}{(1 + \lambda_k)^2} (1 + \lambda_k - B_k^T \tilde{S} B_k) \right)_{k=1}^r.$$

Since $B_k^T P^{-1} B_l = 0$, for $1 \leq k \neq l \leq r$, it follows that $\frac{\partial^2 F_n^i}{\partial \zeta_k \partial \zeta_l}(\theta) = 0$ for $k \neq l$, $i = 1, 2$. Also,

$$\frac{\partial^2 F_n^1}{\partial \zeta_k^2}(\theta) = e^{\zeta_k} B_k^T P^{-1} \tilde{S} P^{-1} B_k \left[2e^{\zeta_k} (B_k^T P^{-1} B_k) - 1 \right] = \frac{\lambda_k (\lambda_k - 1)}{(1 + \lambda_k)^3} B_k^T \tilde{S} B_k,$$

$$\frac{\partial^2 F_n^2}{\partial \zeta_k^2}(\theta) = e^{\zeta_k} B_k^T P^{-1} B_k \left[1 - e^{\zeta_k} (B_k^T P^{-1} B_k) \right] = \frac{\lambda_k}{(1 + \lambda_k)^2}.$$

Thus, the Hessian operator of $\tilde{L}_n(\cdot)$ w.r.t. ζ is given by

$$H_{n,\zeta}(\theta) = \text{diag} \left(\frac{\lambda_k}{(1 + \lambda_k)^3} \left((\lambda_k - 1) B_k^T \tilde{S} B_k + (1 + \lambda_k) \right) \right)_{k=1}^r. \quad (64)$$

Boundedness and inversion of $H(\theta_*; \theta_*)$

As discussed in Section 6, the Hessian operator $H(\theta_*; \theta_*)$ is “block diagonal”. So, we only need to show the boundedness and calculate the inverse of $H_B(\theta_*; \theta_*)$ and $H_{\zeta}(\theta_*; \theta_*)$. First note that, from (64) we have,

$$H_{\zeta}(\theta_*; \theta_*) = \text{diag} \left(\frac{\lambda_{*k}}{(1 + \lambda_{*k})^3} ((\lambda_{*k} - 1) B_{*k}^T \Gamma_* B_{*k} + (1 + \lambda_{*k})) \right)_{k=1}^r = \Lambda_*^2 (I_r + \Lambda_*)^{-2},$$

which is clearly positive definite with eigenvalues bounded away from 0 and ∞ , due to conditions **A1’** and **C’’**.

Next, we show that $H_B(\theta_*; \theta_*)(X, X) \geq C \langle X, X \rangle_c$, for some $C > 0$, for all $X \in \mathcal{T}_{B_*}$. Define $F_B^1(\theta_*; \theta_*) = \mathbb{E}_{\theta_*} F_{n,B}^1(\theta_*)$ and $G_{BB}^1(\theta_*; \theta_*) = \mathbb{E}_{\theta_*} G_{n,BB}^1(\theta_*)$. Note that $H_B(\theta_*; \theta_*)$ is obtained by replacing $F_{n,B}^1(\theta_*)$ and $G_{n,BB}^1(\theta_*)$ by $F_B^1(\theta_*)$ and $G_{BB}^1(\theta_*; \theta_*)$, respectively, in (62). Observe that,

$$F_B^1(\theta_*) = -2\Gamma_* B_* Q_*^{-1} = -2B_*(I_r + \Lambda_*)Q_*^{-1} = -2B_*\Lambda_*,$$

where $Q_* = Q(\theta_*) = \Lambda_*^{-1}(I_r + \Lambda_*)$, and we have used the fact that $\Gamma_* B_* = B_*(I_r + \Lambda_*)$. For notational simplicity we use F_B^1 to denote $F_B^1(\theta_*)$. Note that, for $X \in \mathcal{T}_{B_*}$, $X = B_* A_X + (I - B_* B_*^T) C_X$, where $A_X = -A_X^T \in \mathbb{R}^{r \times r}$, and $C_X \in \mathbb{R}^{M \times r}$. Using this representation, for any $X, Y \in \mathcal{T}_{B_*}$, we have

$$\begin{aligned} \frac{1}{2} \text{tr} [((F_B^1)^T X B_*^T + B_*^T X (F_B^1)^T) Y] &= -\text{tr} [\Lambda_* B_*^T X B_*^T Y + B_*^T X \Lambda_* B_*^T Y] \\ &= \text{tr} [\Lambda_* X^T B_* B_*^T Y + B_*^T X \Lambda_* Y^T B_*] = 2\text{tr} [\Lambda_* X^T B_* B_*^T Y], \end{aligned} \quad (65)$$

and

$$\begin{aligned} -\frac{1}{2} \text{tr} [(B_*^T (F_B^1) + (F_B^1)^T B_*) X^T (I_M - B_* B_*^T) Y] \\ = \text{tr} [(B_*^T B_* \Lambda_* + \Lambda_* B_*^T B_*) X^T (I_M - B_* B_*^T) Y] = 2\text{tr} [\Lambda_* X^T (I_M - B_* B_*^T) Y]. \end{aligned} \quad (66)$$

Next, notice that, for $X \in \mathcal{T}_{B*}$, $G_{BB}^1(\theta_*; \theta_*)(X) = -2\Gamma_* X Q_*^{-1}$. Therefore,

$$\begin{aligned} \text{tr} [Y^T G_{BB}^1(\theta_*; \theta_*)(X)] &= -2\text{tr} [Y^T \Gamma_* X Q_*^{-1}] \\ &= -2\text{tr} [Y^T (I_M - B_* B_*^T) X Q_*^{-1}] - 2\text{tr} [Y^T B_* (I_r + \Lambda) B_*^T X Q_*^{-1}]. \end{aligned} \quad (67)$$

Now, combining (65), (66) and (67), and using the definition of $H_B(\theta_*; \theta_*)$, and the facts that $I_r - Q_*^{-1} = (I_r + \Lambda_*)^{-1}$, $X^T B_* = A_X^T = -A_X$ and $B_*^T Y = A_Y = -A_Y^T$, after some simple algebra we have,

$$\begin{aligned} &H_B(\theta_*; \theta_*)(X, Y) \\ &= 2 \left(\text{tr} [X^T B_* \Lambda_* B_*^T Y (I_r + \Lambda_*)^{-1}] - \text{tr} [X^T B_* B_*^T Y Q_*^{-1}] \right) + 2\text{tr} [\Lambda_*^2 (I_r + \Lambda_*)^{-1} X^T (I_r - B_* B_*^T) Y] \end{aligned} \quad (68)$$

Again, since $A_X = -A_X^T$ and $A_Y = -A_Y^T$, denoting by $A_{X,ij}$ and $A_{Y,ij}$, the (i, j) -th element of A_X and A_Y respectively, we have

$$\begin{aligned} &\text{tr} [X^T B_* \Lambda_* B_*^T Y (I_r + \Lambda_*)^{-1}] - \text{tr} [X^T B_* B_*^T Y Q_*^{-1}] \\ &= -\text{tr} [A_X (\Lambda_* A_Y (I_r + \Lambda_*) - A_Y \Lambda_* (I_r + \Lambda_*)^{-1})] \\ &= -\sum_{i=1}^r \sum_{j=1}^r A_{X,ij} \left(\frac{\lambda_{*j}}{1 + \lambda_{*i}} A_{Y,ji} - \frac{\lambda_{*i}}{1 + \lambda_{*i}} A_{Y,ji} \right) \\ &= \sum_{i=1}^r \sum_{j=1}^r A_{X,ij} A_{Y,ij} \left(\frac{\lambda_{*j} - \lambda_{*i}}{1 + \lambda_{*i}} \right) = \sum_{i=1}^{j-1} \sum_{j=i+1}^r A_{X,ij} A_{Y,ij} \left[(\lambda_{*i} - \lambda_{*j}) \left(\frac{1}{1 + \lambda_{*j}} - \frac{1}{1 + \lambda_{*i}} \right) \right] \\ &= \sum_{i=1}^{j-1} \sum_{j=i+1}^r A_{X,ij} A_{Y,ij} \frac{(\lambda_{*i} - \lambda_{*j})^2}{(1 + \lambda_{*i})(1 + \lambda_{*j})} = \frac{1}{2} \sum_{i=1}^r \sum_{j=1}^r A_{X,ij} A_{Y,ij} \frac{(\lambda_{*i} - \lambda_{*j})^2}{(1 + \lambda_{*i})(1 + \lambda_{*j})}. \end{aligned} \quad (69)$$

Since $\min_{1 \leq k \neq k' \leq r} (\lambda_{*k} - \lambda_{*k'})^2 (1 + \lambda_{*k})^{-2} \geq C_{1*}$, and $\lambda_{*r} \geq C_{*2}$, for some constants $C_{1*}, C_{*2} > 0$ (value depending on \bar{c}_1 and \bar{c}_2 appearing in **A1'**), it follows from (68) and (69) that for $X \in \mathcal{T}_{B*}$,

$$\begin{aligned} H_B(\theta_*; \theta_*)(X, X) &\geq C_{*1} \text{tr} (X^T B_* B_*^T X) + 2C_{*2} \text{tr} (X^T (I_M - B_* B_*^T) X) \\ &\geq C_{*3} \text{tr} (X^T X), \end{aligned}$$

where $C_{*3} = \min\{C_{*1}, 2C_{*2}\}$. This proves that $H_B(\theta_*; \theta_*)(X, X)$ is bounded below in the Euclidean norm and hence in the canonical metric because of the norm equivalence. An upper bound follows similarly.

Proof of Corollary 3 : From (68) and (69), we can derive an explicit expression of $H_B^{-1}(\theta_*; \theta_*)$. Note that $H_B^{-1}(\theta_*; \theta_*)(X)$ is defined as

$$H_B(\theta_*; \theta_*)(H_B^{-1}(\theta_*; \theta_*)(X), Y) = \langle X, Y \rangle_c, \quad \text{for any } Y \in \mathcal{T}_{B*}.$$

Therefore, for $X \in \mathcal{T}_{B*}$ and $A_X = B_*^T X$,

$$H_B^{-1}(\theta_*; \theta_*)(X) = \frac{1}{2} B_* \left(\left(\frac{(1 + \lambda_{*i})(1 + \lambda_{*j})}{(\lambda_{*i} - \lambda_{*j})^2} A_{X,ij} \right) \right) + \frac{1}{2} (I_M - B_* B_*^T) X \Lambda_*^{-2} (I_r + \Lambda_*).$$

Using this, we can now get an explicit expression for $H_B^{-1}(\theta_*; \theta_*)(\nabla_B \tilde{L}_n(\theta_*))$. From (61), we have

$$\begin{aligned} B_*^T \nabla_B \tilde{L}_n(\theta_*) &= 2 \left[Q_*^{-1} B_*^T \tilde{S} B_* - B_*^T \tilde{S} B_* Q_*^{-1} \right] = 2 \left(\left(B_{*i}^T \tilde{S} B_{*j} \left(\frac{\lambda_{*i}}{1 + \lambda_{*i}} - \frac{\lambda_{*j}}{1 + \lambda_{*j}} \right) \right) \right) \\ &= 2 \left(\left(\frac{(\lambda_{*i} - \lambda_{*j})}{(1 + \lambda_{*i})(1 + \lambda_{*j})} B_{*i}^T \tilde{S} B_{*j} \right) \right). \end{aligned}$$

Also,

$$(I_M - B_* B_*^T) \nabla_B \tilde{L}_n(B_*, \Lambda_*) = -2(I_M - B_* B_*^T) \tilde{S} B_* Q_*^{-1}.$$

Thus, it follows that

$$\begin{aligned} -H_B^{-1}(\theta_*; \theta_*) \left(\nabla_B \tilde{L}_n(\theta_*) \right) &= - \left(\left(\frac{1}{(\lambda_{*i} - \lambda_{*j})} B_{*i}^T \tilde{S} B_{*j} \right) \right) + (I_M - B_* B_*^T) \tilde{S} B_* \Lambda_*^{-1} \\ &= - \left[\mathbf{R}_1 \tilde{S} B_{*1} : \cdots : \mathbf{R}_r \tilde{S} B_{*r} \right]. \end{aligned}$$

Appendix E : Gradient and Hessian on product manifolds

In this section, we give a brief outline of the intrinsic geometry associated with the product manifold of two Riemannian manifolds, and as an application we consider the manifold $\mathcal{S}_{M,r} \otimes \mathbb{R}^r$, which is the parameter space for (B, ζ) in our problem.

Product Manifolds

Consider two Riemannian manifolds: \mathcal{M}, \mathcal{N} with metrics g_M and g_N , respectively. The product manifold \mathcal{P} of \mathcal{M}, \mathcal{N} is then defined as:

$$\mathcal{P} := \mathcal{M} \otimes \mathcal{N} = \{(x, y) : x \in \mathcal{M}, y \in \mathcal{N}\}$$

with the tangent space at a point $p = (x, y) \in \mathcal{P}$,

$$\mathcal{T}_p \mathcal{P} := \mathcal{T}_x \mathcal{M} \oplus \mathcal{T}_y \mathcal{N}$$

where $\mathcal{T}_x \mathcal{M}, \mathcal{T}_y \mathcal{N}$ are tangent spaces of \mathcal{M}, \mathcal{N} at points x, y , respectively. The Riemannian metric g on the tangent space $\mathcal{T}\mathcal{P}$ is naturally defined as

$$\langle T_1, T_2 \rangle_g := \langle \xi_1, \xi_2 \rangle_{g_M} + \langle \eta_1, \eta_2 \rangle_{g_N},$$

where $T_i = (\xi_i, \eta_i) \in \mathcal{T}\mathcal{P}$, with $\xi_i \in \mathcal{T}\mathcal{M}$ and $\eta_i \in \mathcal{T}\mathcal{N}$ ($i = 1, 2$).

By the above definition of the product manifold \mathcal{P} , the intrinsic gradient and Hessian of a smooth function f defined on \mathcal{P} are as follows:

- Gradient:

$$\nabla f = (\nabla_{\mathcal{M}} f_{\mathcal{M}}, \nabla_{\mathcal{N}} f_{\mathcal{N}}),$$

where $f_{\mathcal{M}}$ ($f_{\mathcal{N}}$) is f viewed as a function on \mathcal{M} (\mathcal{N}); and $\nabla_{\mathcal{M}}$ ($\nabla_{\mathcal{N}}$) denotes the gradient operator for functions defined on \mathcal{M} (\mathcal{N}).

- Hessian: for $T_i = (\xi_i, \eta_i) \in \mathcal{TP}$ ($i = 1, 2$),

$$\begin{aligned} H_f(T_1, T_2) &= H_{f_{\mathcal{M}}}(\xi_1, \xi_2) + \langle \nabla_{\mathcal{N}} \langle \nabla_{\mathcal{M}} f_{\mathcal{M}}, \xi_1 \rangle_{g_{\mathcal{M}}}, \eta_2 \rangle_{g_{\mathcal{N}}} \\ &+ \langle \nabla_{\mathcal{M}} \langle \nabla_{\mathcal{N}} f_{\mathcal{N}}, \eta_1 \rangle_{g_{\mathcal{N}}}, \xi_2 \rangle_{g_{\mathcal{M}}} + H_{f_{\mathcal{N}}}(\eta_1, \eta_2). \end{aligned}$$

The above expression is derived from the bi-linearity of the Hessian operator and its definition. Also note that

$$\langle \nabla_{\mathcal{N}} \langle \nabla_{\mathcal{M}} f_{\mathcal{M}}, \xi_1 \rangle_{g_{\mathcal{M}}}, \eta_2 \rangle_{g_{\mathcal{N}}} = \langle \nabla_{\mathcal{M}} \langle \nabla_{\mathcal{N}} f_{\mathcal{N}}, \eta_2 \rangle_{g_{\mathcal{N}}}, \xi_1 \rangle_{g_{\mathcal{M}}}.$$

Application to the product of a Stiefel manifold and an Euclidean space

Consider the special case: $\mathcal{M} = \mathcal{S}_{M,r}$ with the canonical metric $\langle \cdot, \cdot \rangle_c$, and $\mathcal{N} = \mathbb{R}^d$ with Euclidean metric. For a point $p = (B, x)$ on the product manifold \mathcal{P} , the tangent space is

$$\mathcal{T}_p \mathcal{P} = \mathcal{T}_B \mathcal{M} \oplus \mathcal{T}_x \mathcal{N},$$

where

$$\mathcal{T}_B \mathcal{M} = \{\Delta \in \mathbb{R}^{M \times r} : B^T \Delta = -\Delta^T B\}, \quad \text{and} \quad \mathcal{T}_x \mathcal{N} = \mathbb{R}^d.$$

For a smooth function f defined on the product space \mathcal{P} :

- Gradient (at p):

$$\nabla f|_p = \left(\nabla_{\mathcal{M}} f, \frac{\partial f}{\partial x} \right) \Big|_p,$$

where $\nabla_{\mathcal{M}} f|_p = f_B - B f_B^T B$ (with $f_B = \frac{\partial f}{\partial B}$).

- Hessian operator (at p): for $T = (\Delta, a)$, and for $X = (X_B, \eta) \in \mathcal{T}_p \mathcal{P}$,

$$H_f(T, X)|_p = H_{f_{\mathcal{M}}}(\Delta, X_B) + \left\langle \frac{\partial}{\partial x} \langle \nabla_{\mathcal{M}} f, \Delta \rangle_c, \eta \right\rangle + \left\langle \frac{\partial}{\partial x} \langle \nabla_{\mathcal{M}} f, X_B \rangle_c, a \right\rangle + a^T \frac{\partial^2 f}{\partial x^2} \eta, \quad (70)$$

where

$$H_{f_{\mathcal{M}}}(\Delta, X_B)|_p = f_{BB}(\Delta, X_B) + \frac{1}{2} \text{Tr} [(f_B^T \Delta B^T + B^T \Delta f_B^T) X_B] - \frac{1}{2} \text{Tr} [(B^T f_B + f_B^T B) \Delta^T \Pi X_B],$$

with $\Pi = I - BB^T$.

- Inverse of Hessian operator (at p): for $G \in \mathcal{T}_p \mathcal{P}$, $T = H_f^{-1}(G)|_p$ is defined as: $T = (\Delta, a) \in \mathcal{T}_p \mathcal{P}$ such that for any $X = (X_B, \eta) \in \mathcal{T}_p \mathcal{P}$ the following equation is satisfied

$$H_f(T, X)|_p = \langle G, X \rangle_g \quad .$$

Appendix F : Some inequalities involving matrices

In this paper we make frequent use of the following matrix inequalities:

- For any A, B ,

$$\|AB\|_F \leq \|A\|_F \|B\|, \quad \text{and} \quad \|AB\|_F \geq \|A\|_F \lambda_{\min}(B), \quad (\text{for } B \text{ positive definite})$$

where $\lambda_{\min}(B)$ is the smallest eigenvalue of B . Also, if A and B are invertible then

$$A^{-1} - B^{-1} = A^{-1}(B - A)B^{-1} = B^{-1}(B - A)A^{-1}. \quad (71)$$

- *Weilandt's inequality (Horn and Johnson (1994))*: For symmetric $p \times p$ matrices A, B with eigenvalue sequences $\lambda_1(A) \geq \dots \geq \lambda_p(A)$ and $\lambda_1(B) \geq \dots \geq \lambda_p(B)$, respectively,

$$\sum_{i=1}^p |\lambda_i(A) - \lambda_i(B)|^2 \leq \|A - B\|_F^2 \quad (72)$$

- *Eigenvector perturbation (Paul (2005))*: Let A be a $p \times p$ positive semidefinite matrix, with j -th largest eigenvalue $\lambda_j(A)$ with corresponding eigenvector \mathbf{p}_j , and $\tau_j := \max\{(\lambda_{j-1}(A) - \lambda_j(A))^{-1}, (\lambda_j(A) - \lambda_{j+1}(A))^{-1}\}$ is bounded (we take $\lambda_0(A) = \infty$ and $\lambda_{p+1}(A) = 0$). Let B be a symmetric matrix. If \mathbf{q}_j denotes the eigenvector of $A + B$ corresponding to the j -th largest eigenvalue (which is of multiplicity 1, for $\|B\|$ small enough, by (72)), then (assuming without loss of generality $\mathbf{q}_j^T \mathbf{p}_j > 0$),

$$\|\mathbf{q}_j - \mathbf{p}_j\| \leq 5 \frac{\|B\|}{\tau_j} + 4 \left(\frac{\|B\|}{\tau_j} \right)^2. \quad (73)$$

Reference

1. Antoniadis, A. and Sapatinas, T. (2007) : Estimation and inference in functional mixed-effects models. *Computational Statistics and Data Analysis* **51**, 4793-4813.
2. Bickel, P. J. and Levina, E. (2007) : Covariance regularization by thresholding. *Technical report #744*, Department of Statistics, UC Berkeley. *Annals of Statistics*. To appear.
3. Bickel, P. J. and Levina, E. (2008) : Regularized estimation of large covariance matrices. *Annals of Statistics* **36**, 199-227.
4. Burman, P. (1985) : A data dependent approach to density estimation. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* **69**, 609-628.
5. Chen, A. and Bickel, P. J. (2006) : Efficient independent component analysis. *Annals of Statistics* **34**, 2825-2855.
6. de Boor, C. (1974) : Bounding the error in spline interpolation. *SIAM Review* **16**, 531-544.
7. de Boor, C. (1978) : *A Practical Guide to Splines*. Springer-Verlag.
8. DeVore, R. A. and Lorentz, G. G. (1993) : *Constructive Approximation*. Springer.

9. El Karoui, N. (2007) : Operator norm consistent estimation of large dimensional sparse covariance matrices. *Annals of Statistics*. To appear.
10. Davidson, K. R., and Szarek, S. (2001) : Local operator theory, random matrices and Banach spaces. In “*Handbook on the Geometry of Banach spaces*”, (V. 1, Johnson, W. B., Lindenstrauss, J. eds.), 317-366, Elsevier Science.
11. Edelman, A., Arias, T. A. and Smith, S. T. (1998) : The geometry of algorithms with orthogonality constraints, *SIAM Journal on Matrix Analysis and Applications* **20**, 303-353.
12. Fan, J., Fan, Y. and Lv, J. (2006) : High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics*, (to appear).
13. Ferraty, F. and Vieu, P. (2006) : *Nonparametric Functional Data Analysis : Theory and Practice*. Springer.
14. Hall, P. and Hosseini-Nasab, M. (2006) : On properties of functional principal components analysis. *Journal of the Royal Statistical Society, Series B* **68**, 109-126.
15. Hall, P., Müller, H.-G. and Wang, J.-L. (2006) Properties of principal component methods for functional and longitudinal data analysis. *Annals of Statistics* **34**, 1493-1517.
16. Horn, R. A. and Johnson, C. R. (1994) : *Topics in Matrix Analysis*. Cambridge University Press.
17. James, G. M., Hastie, T. J. and Sugar, C. A. (2000) : Principal component models for sparse functional data. *Biometrika*, **87**, 587-602.
18. Kato, T. (1980) : *Perturbation Theory of Linear Operators*. Springer-Verlag.
19. Muirhead, R. J. (1982) : *Aspects of Multivariate Statistical Theory*, John Wiley & Sons.
20. Oller, J. M. and Corcuera, J. M. (1995) : Intrinsic analysis of statistical estimation. *Annals of Statistics* **23**, 1562-1581.
21. Paul, D. (2005) : *Nonparametric Estimation of Principal Components*. Ph. D. Thesis. Stanford University.
22. Paul, D. and Johnstone, I. M. (2007) : Augmented sparse principal component analysis for high dimensional data. *Working Paper*.
(<http://anson.ucdavis.edu/~debashis/techrep/augmented-spca.pdf>)
23. Peng, J. and Paul, D. (2007) : A geometric approach to maximum likelihood estimation of covariance kernel from sparse irregular longitudinal data. *Technical Report*. arXiv:0710.5343v1 [stat.ME]. (also at <http://anson.ucdavis.edu/~jie/pd-cov-likelihood-technical.pdf>)
24. Ramsay, J. and Silverman, B. W. (2005) : *Functional Data Analysis, 2nd Edition*. Springer.
25. Yang, A. and Barron, A. (1999) : Information-theoretic determination of minimax rates of convergence. *Annals of Statistics* **27**, 1564-1599.
26. Yao, F., Müller, H.-G. and Wang, J.-L. (2005) : Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association* **100**, 577-590.

D. Paul
Department of Statistics
University of California
Davis, CA 95616
debashis@wald.ucdavis.edu

J. Peng
Department of Statistics
University of California
Davis, CA 95616
jie@wald.ucdavis.edu